AUGUST 1989

COMPUTER

Scientific Visualization

Bringing data into focus

IEEE COMPUTER SOCIETY

THE INSTITUTE OF ELECTRICAL AND
ELECTRONICS ENGINEERS, INC

Advertisement

Parallel Processing

with the Inmos Transputer

from Computer System Architects

APPLICATIONS

- **High-Performance · Cost Effective**

  For computationally intensive applications such as simulation, signal processing and graphics rendering. For image processing and pattern recognition. For distributed databases. For real-time applications like transaction processing, data acquisition and robotics. Using from 1 to 64 (or more) processors. Incrementally expandable and field upgradable.

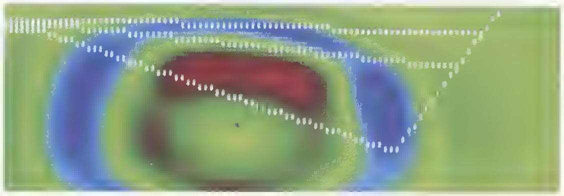

Earthquake simulation of Salt Lake Valley courtesy of University of Utah

- **Acceleration for PC's and Workstations**

  For applications requiring a large, flat, memory address space. For applications which may benefit from incrementally scalable performance by simply adding processors. Put from 1 to 16 processors (with up to 32MB) inside the PC or use just 1 slot to escape to an external chassis. One transputer is up to 3 times faster than a Deskpro 386/25 with 80387!

HARDWARE

- **Compute Nodes**

  Processors: Inmos T425 or T800 32-bit transputers. Up to 12.5 MIPS and 1.25 64-bit MFLOPS (with no vector pipelining required). Two on-chip timers facilitate real-time applications. Hardware scheduler supports multiple processes on each processor. **Links:** Each processor has 4 bidirectional serial communication links, providing data transfer between nodes at up to 9MB/sec (overlapped with computation). Links are differentially buffered and user-configurable, allowing multi-cabinet systems. **Memory:** Each node has 4KB of 50 nsec SRAM and up to 8MB of local DRAM. **Size:** 1 to 4 nodes per board, with up to 20 boards per cabinet. Boards also fit in PC's and AT's. **Cost:** Compute nodes range in price from $900 (a 20 MHz T425 with 256KB) to $5,995 (a 25MHz T800 with 8MB).

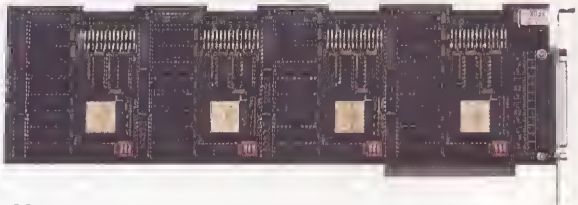

CSA PART 6A: Four T800 transputers with 1 MByte RAM per processor

- **High-Speed Interfaces and Accessories**

  Interfaces for PC and AT hosts. Also for Apollo and SUN-386i workstations which use the AT bus. For hosts and peripherals using the VME bus, including the Apollo-10000 and SUN-3 and -4. A SCSI interface is even available, allowing direct connection of disks and other SCSI peripherals. Also for hosts with a SCSI port, like the Macintosh.

  Accessories include a transputer-based **color graphics controller,** a 32 x 32 **link switch** (allowing programmable reconfiguration of link connections), and a **proto-board** (with 16-bit T222 processor and SRAM/EPROM) facilitating custom data acquisition and real-time control.

SOFTWARE

- **ANSI Standard C · FORTRAN 77 · Modula 2 · Occam 2 · Pascal**

  Compilers come with transputer network loaders, concurrency and message passing libraries, and host server programs for DOS or UNIX. Prices start at $400 (a complete C development system).

- **Operating Environments and Libraries**

  Express from ParaSoft and **Helios** from Perihelion. Debuggers. Math libraries from Topexpress. Graphics libraries. And more.

CSA

Computer System Architects
950 N. University Avenue
Provo, Utah 84604
(801) 374-2300
FAX: 374-2306

Call, write or FAX for more information.

Bringing the Power of the Transputer to Science and Industry

Reader Service Number 1

Breakthrough in presentation of simulation results
SIMSCRIPT II.5 with SIMGRAPHICS
Now you see an animated picture of the system

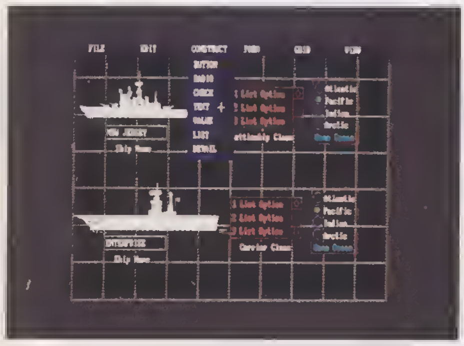

SIMGRAPHICS™ menu builder

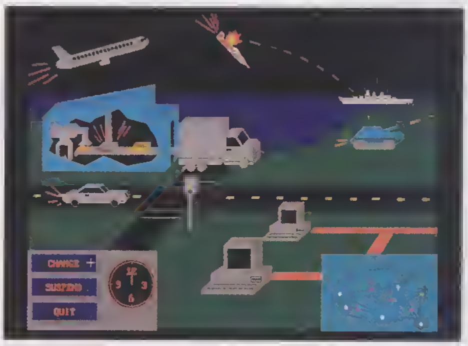

Draw your own icons--or use ours

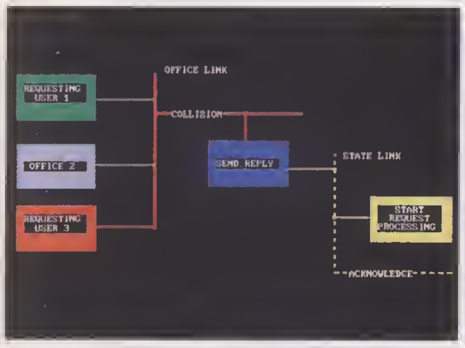

Local Area Network--NETWORK II.5®

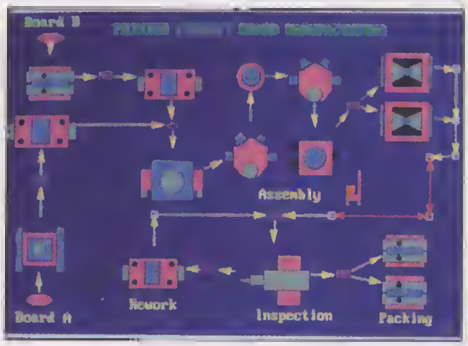

Factory in operation--SIMFACTORY II.5®

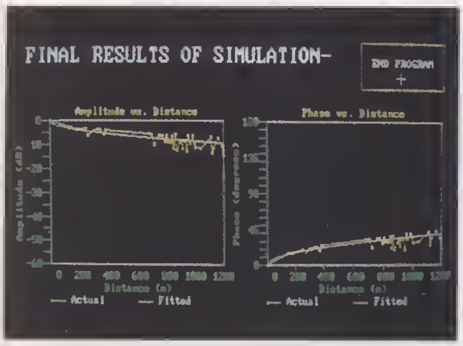

Presentation graphics

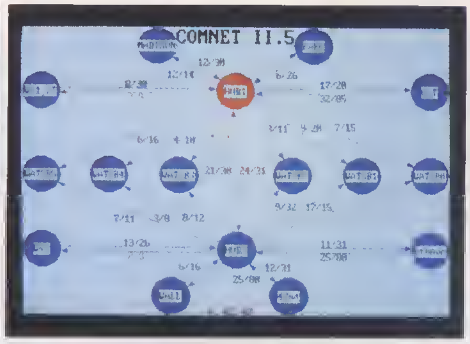

Communications--COMNET II.5®

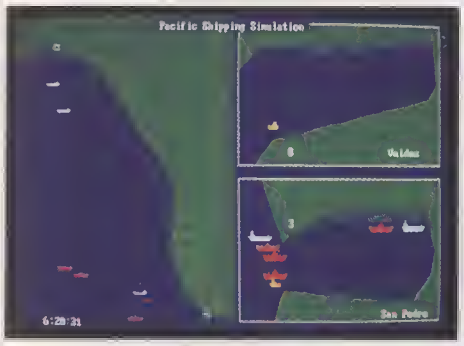

Transportation system

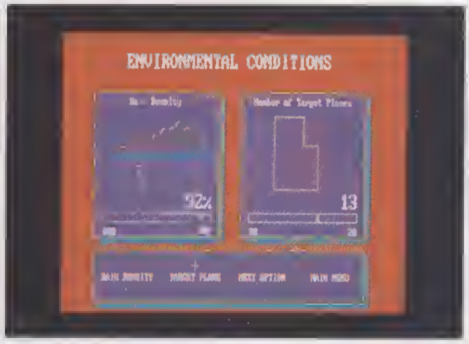

SIMGRAPHICS advanced user interface

Free trial and training

See for yourself how simulation results are now easier to understand.

The free trial contains everything you need to try SIMGRAPHICS™ on your computer.

We send you SIMSCRIPT II.5, animated models, and complete documentation. You can build your own model or modify one of ours.

Try the SIMSCRIPT II.5® language, the timeliness of our support, the accuracy of our documentation, and the facilities for error checking--**everything you need for a successful project.**

For 26 years CACI has provided trial use of its simulation software--**no cost, no obligation.**

Act now for free training

For a limited time we will also include free training.

For immediate information

Call Hal Duncan at (619) 457-9681, FAX (619) 457-1184. In Europe, call Richard Eve on (01) 528-7980, FAX (01) 528-7988.

Rush information on SIMSCRIPT II.5 with SIMGRAPHICS

Limited offer-return the coupon today and we will also include one free course enrollment worth $950.

☐ Send information on your Special University Offer.

Name

Organization

Address

City State Zip

Telephone

Computer Operating System

IEEE COMP

Return to:
CACI Products Company
3344 North Torrey Pines Court
La Jolla, California 92037
Call Hal Duncan at (619) 457-9681.
FAX (619) 457-1184.

In Europe:
CACI Products Division
Regent House, 89 Kingsway
London, WC2B 6RH, United Kingdom
Call Richard Eve on (01) 528-7980.
FAX (01) 528-7988.

SIMSCRIPT II.5, NETWORK II.5, SIMFACTORY II.5, and SIMGRAPHICS are registered trademarks and service marks of CACI, INC. COMNET II.5 is a trademark and service mark of CACI, INC. ©1989 CACI, INC.

COMPUTER

August 1989 Published by the IEEE Computer Society Vol. 22, No. 8

ARTICLES

10 **Guest Editor's Introduction: Visualization in Scientific Computing**
Gregory M. Nielson

12 **Visualization: Expanding Scientific and Engineering Research Opportunities**
Thomas A. DeFanti, Maxine D. Brown, and Bruce H. McCormick

Visualization holds great promise for computational science and engineering, provided we can meet the immediate and long-term needs of both toolmakers and tool users.

27 **Representation and Display of Vector Field Topology in Fluid Flow Data Sets**
James Helman and Lambertus Hesselink

A representation of the global topology based on an analysis of critical points and their connections is much more readily manipulated and visualized than the original data set.

39 **Acquisition and Representation of 2D and 3D Data from Turbulent Flows and Flames**
Marshall B. Long, Kevin Lyons, and Joseph K. Lam

Laser light-scattering techniques for making two- and three-dimensional measurements and better methods for representing the large quantities of data obtained may help solve the "turbulence problem."

46 **Interactive Visualization of 3D Medical Data**
Henry Fuchs, Marc Levoy, and Stephen M. Pizer

Though several techniques for displaying multidimensional data are being developed, research suggests that real-time volume rendering will be preferred for many clinical applications.

53 **Visualizing Large Data Sets in the Earth Sciences**
William Hibbard and David Santek

McIDAS, an interactive visualization system, is revolutionizing the ability of earth scientists to manage and analyze data from remote sensing instruments and numerical simulation models.

60 **The Role of Visualization in the Simulation of Quantum Electronic Transport in Semiconductors**
Norman C. Kluksdahl, Alfred M. Kriman, and David K. Ferry

Using the Wigner function to simulate semiconductor processes, then depicting the results graphically, can lead to new insights concerning these devices.

68 **Scientific Visualization at Research Laboratories**
Section Editor: Lawrence J. Rosenblum

Large scientific research centers are developing visualization methodologies and tools. This special section examines work at Los Alamos, Lawrence Livermore, the National Center for Supercomputing Applications, NASA Ames, and the Naval Research Laboratory.

Circulation: *Computer* (ISSN 0018-9162) is published monthly by the IEEE Computer Society, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720, phone (714) 821-8380. IEEE Computer Society Headquarters, 1730 Massachusetts Ave. NW, Washington, DC 20036-1903; IEEE Headquarters, 345 East 47th St., New York, NY 10017. Annual subscription included in society member dues. Nonmember subscription rates: available upon request. Single copy prices: members $10.00; nonmembers $20.00. This magazine is also available in microfiche form.

Postmaster: Send undelivered copies and address changes to *Computer*, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. Second class postage is paid at New York, New York, and at additional mailing offices.

Copyright and reprint permission: Copyright © 1989 by the Institute of Electrical and Electronics Engineers, Inc. All rights reserved. Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limits of US copyright law for private use of patrons: (1) those post-1977 articles that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 29 Congress Street, Salem, MA 01970; (2) pre-1978 articles without fee. Instructors are permitted to photocopy isolated articles for noncommercial classroom use without fee. For other copying, reprint, or republication permission, write to Permissions Editor, *Computer*, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720.

Editorial: Unless otherwise stated, bylined articles, as well as product and service descriptions, reflect the author's or firm's opinion. Inclusion in *Computer* does not necessarily constitute endorsement by the IEEE or the Computer Society. All submissions are subject to editing for style, clarity, and space considerations. *Computer* subscribes to the Computer Press Association's code of professional ethics.

DEPARTMENTS

4 President's Message
8 Letters to the Editor
104 Open Channel
105 Standards
108 Computer Society News
109 Update
111 Product Reviews
123 New Products
129 IC Announcements
130 Microsystem Announcements
132 Conferences
134 Calendar
140 Call for Papers
147 Book Reviews
151 Contents: CS Magazines
151 New Literature

Computer Society Information Page, inside back cover; Membership Application, 152; Career Opportunities, 144; Advertising and Product Indexes, 128; Reader Service Card, 128A

Scientific visualization makes the unseeable seeable. This example, Romboy Homotopy, was created with software developed on the Cray SMP/48 by Ray Idaszak, Donna Cox, and George Francis of the National Center for Supercomputing Applications. (They also used an Alliant FX/80 and a Silicon Graphics IRIS 4D/70G.) The Romboy Homotopy is a complex deformation in 4D space. The parameterization was created by Francois Apery in 1984 and extended by Francis. This was the first visualization of the homotopy, and it resulted in the discovery of a new surface. The unique shading algorithm used to render these surfaces can correctly smooth-shade single-sided surfaces, as depicted here.

Image, ©1988 Ray Idaszak, Donna Cox, and George Francis, first shown at SIGGraph 88

Cover design by Jay Simpson, Design & Direction

COMPUTER
10662 Los Vaqueros Circle, Los Alamitos, CA 90720

Editor-in-Chief
Bruce D. Shriver, University of Southwestern Louisiana

Editorial Board
Dharma P. Agrawal, North Carolina State University
Dennis R. Allison, Stanford University
Fletcher Buckley, GE Aerospace
Jon T. Butler, Naval Postgraduate School
Scott Davidson, AT&T
Richard Eckhouse, MOCO
Scott E. Fahlman, Carnegie Mellon University
Edmund Gallizzi, Eckerd College
John R. Gurd, University of Manchester
Brent T. Hailpern, IBM T.J. Watson Research Center
James H. Haynes, University of California, Santa Cruz
Guy Johnson, Rochester Institute of Technology
Yale Patt, University of Michigan
Guylaine M. Pollock, Sandia National Laboratories
William Riddle, Software Design & Analysis
Martha Sloan, Michigan Technological University
Akinori Yonezawa, Tokyo Institute of Technology

Submissions
Send six double-spaced copies of articles and special-issue proposals to Bruce D. Shriver, Vice President for Research, University of Southwestern Louisiana, Drawer 42730, Lafayette, LA 70504, phone (318) 231-6000. For electronic submission, Shriver's addresses are b.shriver on Compmail + and shriver@usl.edu on Internet.

Staff
Editor and Publisher: True Seaborn
Assistant Publisher: Douglas L. Combs
Assistant to the Publisher: Pat Paulsen

Managing Editor: Marilyn Potes
Staff Editors: Nancy Hays, Chuck Governale
Steve H. Wilcox
Contributing Editor: Ware Myers

Art Director: Jay Simpson
Article Layout/Design: Vic Grenrock
Membership/Circulation: Christina Champion
Advertising Coordinators: Heidi Rex
Marian Tibayan

Magazine Advisory Committee
Sushil Jajodia (chair), Jon T. Butler, Sumit DasGupta, Joe Hootman, Ted Lewis, David Pessel, True Seaborn, Bruce D. Shriver, John Staudhammer

Publications Board
Sallie Sheppard (chair), James H. Aylor, Victor Basili, P. Bruce Berra, J. Richard Burke, J.T. Cain, David Choy, James Cross, Sumit DasGupta, James J. Farrell III, Joe Hootman, Anil K. Jain, Glen G. Langdon, Ted Lewis, Ming T. Liu, David Pessel, C.V. Ramamoorthy, Bruce D. Shriver, John Staudhammer, Harold Stone, Steven L. Tanimoto

Computer Society President's MESSAGE

Conferences and tutorials

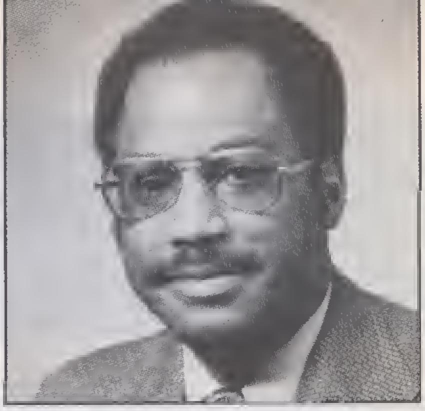

Kenneth R. Anderson, president

Joseph E. Urban, first vice president

Six major conferences we sponsored/cosponsored in 1988 were the Design Automation Conference, the International Test Conference, Supercomputing, the International Conference on Artificial Intelligence Applications, the International Conference on Software Engineering, and the International Conference on Computer-Aided Design. Over the years, interests have shifted, and the major conferences are now ones that focus on the key technologies and disciplines that will drive computer technology in the 1990s.

Sponsoring conferences and tutorials comprises approximately 40 percent of our current program, and it ranks second, behind publications, in generating income. It is a direct result of the activities of the Technical Activities Board and our 33 technical committees. The management of our conference program has been delegated to a group of active volunteers on the Conferences and Tutorials Board. The board is presided over by First Vice President Joseph Urban, who describes this important program in the following message.

Kenneth R. Anderson

One of the technical information aspects of the Computer Society mission, as stated in the constitution, is to "hold meetings for the presentation and discussion of technical papers." To this end, the Computer Society currently sponsors, cosponsors, or cooperates in approximately 100 technical meetings per year. At first glance that appears to be about one technical meeting per thousand members and one meeting every three or four days. In fact, these technical meetings are open to both members and nonmembers (with a discounted attendance fee for members), have anywhere from 50 to 5,000 attendees, and occur throughout the year (with an explicit effort to spread the dates in any particular functional area). The meetings, many of which are sponsored by technical committees, take the form of conferences, symposia, and workshops and often include half-day or full-day tutorials. The Computer Society Conferences and Tutorials Board oversees these technical meetings and tutorials.

Simply holding meetings would not satisfy the Computer Society mission, which specifies that they "maintain a high professional standard." To this end, Mark Haas, the board's planning/self-assessment chair, has been building on the earlier work of Duncan Lawrie to help both the board and the meeting organizers set quality benchmarks. If conference paper citations and meeting attendance numbers are viewed as benchmarks, then the society's technical meetings are in a strong position.

Coverage across the breadth of functional areas and locations are the concern of Boumediene Belkhouche, who chairs the board's New Thrust Areas Committee. This work helps find where the society can fill a niche with a technical meeting. It also helps identify the entrepreneurial conference organizers who form the spirit and effort in conducting these meetings.

The Conferences and Tutorials Board is supported administratively by Vice Chair Pei Hsia, Secretary/Finance Chair Doris Carver, and Meetings Representatives Ray Oberly and Chuck Radke. During the past year, the board has focused on improving policies and procedures to obtain smooth, effective operation of technical meetings. These changes have been conference-organizer driven as well as coming through proactive review of procedures.

Another area being investigated this year is expansion of the tutorials program as an effective means for continuing education. Under the leadership of Guylaine Pollock, who chairs the board's Tutorials Committee, society members will have an active part in shaping this expansion.

Operationally, technical meetings are supported by society staff directed by Anne Marie Kelly. The staff members are there to help, and recent changes in staff organization and conference software have been designed to enhance the support given to technical meeting organizers.

There is a role for everyone in the conferences and tutorials area — from attending a conference/tutorial to serving on a conference organizing/program committee. The recently expanded "Calendar" section of *Computer* aids in identifying meetings to attend, as do the pre- and post-conference reports in the "Conferences" department edited by Edmund L. Gallizzi.

I recommend that you attend a technical meeting to help maintain technical vitality and to meet other computer science and engineering professionals. You can decide for yourself which of the two adages applies — all the action is "in the sessions" or as some say "in the hallways." Discuss your professional abilities with the meeting organizers if you want a role with the technical meeting.

The Conferences and Tutorials Board gratefully acknowledges the efforts expended to satisfy the Computer Society mission. This appreciation includes the efforts of attendees, authors who submit papers, reviewers, keynote speakers, panelists, session chairs, tutorial presenters, program committee members, program/general chairs, and those who serve in the many other areas that make technical meetings happen.

If you have comments or suggestions regarding conferences and tutorials, please send me a note.

Joseph E. Urban
Arizona State University
College of Engineering and Applied Sciences
Computer Science Department
Tempe, AZ 85287

Take Our Course In C And The First Lesson You'll Learn Is In Economics.

CLASSROOM INSTRUCTION $1,500
COMPLETE C VIDEO COURSE $ 295
YOU SAVE $1,205

C's power, and portability make it the language of choice for software developers.

Unfortunately, learning C can be a very costly proposition. Classroom instruction is, in a word, expensive. And many C video courses carry hefty price tags.

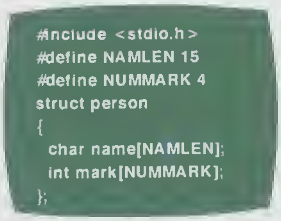
```
#include <stdio.h>
#define NAMLEN 15
#define NUMMARK 4
struct person
{
  char name[NAMLEN];
  int mark[NUMMARK];
};
```

The top C video course at the lowest possible price

But now, there's The Complete C Video Course from Zortech. It's the ultimate C training tool for home or work. And all it costs is $295.

You get ten videos with 36 lessons covering all levels of programming skill. A comprehensive, easy-to-follow 365 page workbook. And even a free C compiler.

Free C compiler included

Yes, that's right. The Complete C Video Course includes our famous C compiler (it runs on any MS-DOS machine) with linker, library manager, full graphics library and on-line help. It's the choice of professional programmers everywhere for fast code, fast development and fast debugging.

Learn C in as little as two weeks

Speaking of speedy, with The Complete C Video Course you can learn C in only two weeks. Compare that with the up to four months it can take to learn C in class.

Each lesson averages 17 minutes of clear, concise instructions. Used in conjunction with our workbook you'll find they provide everything you need to know to become proficient in programming in C.

Save your company thousands

If you think The Complete C Video Course is a great way for you to save money learning C, think about how much it could save your company. Use it instead of sending programmers to school and you'll save thousands. What's more, The Complete C Video Course is even tax deductible.

C is unquestionably the most valuable programming language you can master. And now you can get everything you need to become productive in it from course to compiler to tools for an economical $295. Mail the coupon or call our hotline to receive it ASAP.

ZORTECH

Look at all these C video pluses

- ★ Only $295 complete.
- ★ Ten videos with 36 lessons.
- ★ Comprehensive 365-page workbook.
- ★ Free C compiler with linker, library manager, full graphics library and on-line help.
- ★ Compiler and hardware independent.
- ★ Designed to help you learn C in as little as two weeks.
- ★ Tax deductible.

THE BRITISH PERSONAL COMPUTER AWARDS 1988
PERSONAL COMPUTER WORLD EDITOR'S AWARD
WINNER

Zortech Inc.
366 Massachusetts Ave.
Arlington, MA, 02174
Tel: (617) 646-6703
Fax: (617) 643-7969

★ Yes, rush me The Complete C Video Course including free C compiler for $295.00 (VHS only)

★ Please include (No.   ) extra workbooks at $29.95 each.

★ I'd like to order (No.   ) extra C compilers with this course at the special price of $49.95.

Name/Company . . . .

Address . . . .

. . . .

Phone . . . .

City . . . .

State . . . . Zip . . . .

Here's my check for . . . .

VISA/MC # . . . .

Exp. Date . . . .

%8 *Prices do not include shipping*

The Complete C Video Course $295

Order Hotline (800)848-8408

Reader Service Number 2

Clean Elegant Powerful...

. . . Eiffel

The Object-Oriented Language for Today and Tomorrow

So you're looking for reliability, reusability, and maintainability – so you know you need an object-oriented language – and you want the very best. A language which will stand the test of time – ask around, you'll hear Eiffel over and over again. And for very good reason.

Eiffel is the industrial application of modern software engineering techniques. It offers the full realm of an advanced object-oriented language and environment: single and multiple inheritance, dynamic binding, static typing, genericity and more within a simple, easy-to-learn language.

Eiffel is the only environment which includes reliable software construction techniques: assertions, invariants, rigorous exception handling. Well defined interfaces for your existing software. Robust libraries of reusable components including X-based graphics. A powerful CASE environment ensuring automatic recompilation, computer-aided documentation, graphical browsing of system structures. Portability on all UNIX systems among others. Cross-development via the generation of self-contained C packages. Easy interface with packages and software written in other languages. And incredible support.

Don't worry, you won't be alone. The list of major companies using Eiffel is impressive: Boeing, Philips GTE, British Telecom, Thomson, HP, Sun Microsystems, Cognos, Lawrence Livermore, Tektronix, Sandia, Telecom Australia, BNR, EDF, and on and on.

Eiffel can solve your problems both today and tomorrow. Need more information? Just give us a call. So what are you waiting for?

THE BOOK
Object-Oriented Software Construction, Bertrand Meyer, called a "tour de force" by *IEEE Software*. Order from Prentice-Hall (ISBN 013-629049-3), or directly from Interactive.

THE SEMINAR
Object-Oriented Design and Programming: created and presented by the designers of Eiffel. For information and registration, call us.

SEE EIFFEL
at all the major shows. Call for location and dates.

Interactive
Software Engineering Inc.

270 Storke Road, Suite 7, Goleta, CA 93117 USA
Telephone: (805) 685-1006 FAX: (805) 685-6869

European Distributors: S.O.L., 4, rue René Barthélémy, 92120 Montrouge France Tel. (33) 1 46 57 13 36, FAX (33) 1 46 57 01 03;
ETNOTEAM, Via Staro 4, 20134 Milano, Italy Tel. (39) 2 2141521, FAX (39) 2 2142231

Eiffel is a trademark of Interactive Software Engineering Inc.

Reader Service Number 3

SCIENTIFIC COMPUTING

New in 1989

IMPACT of Computing in Science and Engineering

MANAGING EDITOR
Peter Deuflhard
Konrad Zuse Center
Berlin
Federal Republic of Germany

- Focuses on articles from the areas of mathematical and scientific modeling, scientific computing, computer science, and scientific and engineering applications
- Features interdisciplinary papers touching upon a combination of at least two of the four main areas
- Describes real-life problems attacked by a new or specifically adapted algorithm or by any other new computing technique

Volume 1 (1989), 4 issues
ISSN 0899-8248
In the U.S.A. and Canada: $80.00
All other countries: $92.00

Computer Vision, Graphics, and Image Processing

EDITORS
Linda G. Shapiro
University of Washington, Seattle
Herbert Freeman
Rutgers University
New Brunswick, New Jersey
Thomas S. Huang
Coordinated Science Laboratory
University of Illinois, Urbana
Norman Badler
University of Pennsylvania
Philadelphia
Azriel Rosenfeld
Center for Automation Research
University of Maryland
College Park

- Publishes papers of high quality on the computer processing of pictorial information
- Provides a focal point for literature ordinarily scattered throughout a wide range of journals in computer science, electrical engineering, optics, and other subjects

Volumes 45–48 (1989), 12 issues
ISSN 0734-189X
In the U.S.A. and Canada: $388.00
All other countries: $474.00

Sample copies and privileged personal rates are available upon request.
For more information, please write or call:

ACADEMIC PRESS, INC.
Journal Promotion Department
1250 Sixth Avenue, San Diego, CA 92101
U.S.A.
(619) 699-6742

All prices are in U.S. dollars and are subject to change without notice.

S9183

Reader Service Number 4

LETTERS

Trillium similar to InterViews

To the Editor:

I am writing in response to the excellent article by Linton, Vlissides, and Calder, "Composing User Interfaces with InterViews," published in the February 1989 issue of *Computer* (pp. 8-22).

For several years I have been involved in supporting Trillium, a user interface design environment used by the Xerox Development and Manufacturing Group in the design of new products.*

The groups' experiences with Trillium support most of the points Linton et al. made about the value of separating interaction and abstraction and providing simple ways of composing complex objects and behaviors from primitives. In fact, many of the abstraction and composition mechanisms provided by InterViews are directly analogous to operations and facilities supported by Trillium. However, two points are worth mentioning.

First, there is one major difference between these environments. InterViews is designed to "help programmers create . . . interfaces," and to this end it appears very promising. However, most programmers know little about the ergonomic, operability, and human cognitive factors that must be considered when designing human/machine interfaces. Hence, tools designed primarily to help programmers build user interfaces do not guarantee better interfaces. Trillium, on the other hand, was designed to be used by user interface designers. One objective is reducing the time needed to translate user interface design specifications (often consisting of hand or computer-drawn graphics and hand-written behavioral rules) into code written by software engineers. In addition, Trillium attempts a computer-aided design environment that gives more control to those who know most about human factors and operability.

Second, our experience with Trillium shows that separating presentation from behavioral specification, and providing graphical and behavioral abstractions, is an enormous help in designing prototype user interfaces for computer-controlled systems (supporting Linton et al.). But by far the most difficult aspect of user interface design is behavioral specification. Specifying the behavior of even a very simple object such as an on/off button is programming a machine. Environments like InterViews and Trillium do make programming easier, but it is still programming. A harder problem is providing abstractions that make machine behavior specification easy for designers with primary experience and training in human factors and cognitive psychology. Although InterViews, and the other toolkits mentioned, provide valuable support for programmers building user interfaces, what's really needed is a strong research effort directed at providing user interface tools that support designers, not programmers.

William L. Anderson
Xerox Corp.

---

*D. Austin Henderson, Jr., "The Trillium User Interface Design Environment," *Proc. CHI 86 Human Factors in Computing Systems*, ACM, New York, 1986, pp. 221-227.

How to meet your recommended daily Ada requirements on all Intel-based systems.

Now, the leader in Ada Compilers in the 286/386 DOS world offers a compatible family for all Intel-based systems.

Whatever chip: 80386; 80286; 80186; 8086. Whatever operating system: DOS 16-bit with Alsys Extender; DOS 32-bit with Phar Lap; OS/2; 386/ix, AIX PS/2, SunOS, XENIX, and other UNIX variants; bare 80x86 applications; and, soon, BTOS/CTOS targets. Whatever host platform: PC AT; PS/2; Z-248; 386-based PS/2; Compaq 386; and Sun386i, to name a few.

Whatever you use, it can be fortified with the power and flexibility of an Alsys Ada compiler.

Just complete the coupon or call 617-270-0030 for more information on Alsys Ada Compilation Systems. Because to get everything you require from your Intel-based system, you need a compiler that's technology-enriched.

alsys

Check boxes that interest you.
Please send me information on Alsys Ada compilers for: ☐ 286 DOS ☐ 386 DOS ☐ 386 UNIX ☐ OS/2
Please send me more information on:
☐ Ada Education ☐ Other Alsys Ada compilers
Name __________
Company __________
Address __________
City _____ State ___ Zip _____
Phone (____) __________
Mail to: Alsys, Inc., 67 South Bedford St., Burlington, MA 01803-5152 IEEEC-8/89

AdaNow.

In the US: Alsys Inc., 67 South Bedford St., Burlington, MA 01803-5152, Tel: (617) 270-0030; In the UK: Alsys, Partridge House, Newtown Road, Henley-on-Thames, Oxon RG9 1EN, Tel: 44(491)579090; In France: Alsys, 20 Avenue de Versailles, 78170 La Celle St. Cloud, France, Tel: 33(1) 3918.12.44

Guest Editor's Introduction

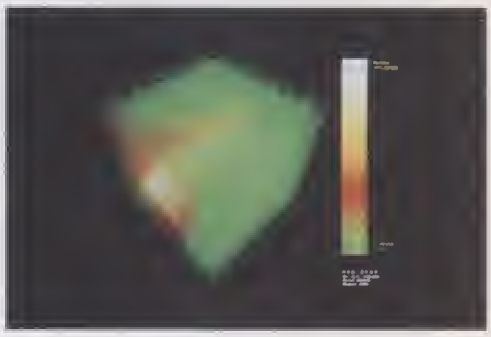

Visualization in

Gregory M. Nielson,

The purpose of computing is insight, not numbers.

— *Richard Hamming*[1]

We all know how a graph can help us understand a relationship. We learn to make and read graphs early in our education, and we reap the benefits almost continuously in our daily lives. The graph of a function $y = f(x)$ or the bar chart for a table of values (x_i, y_i), $i = 1, \ldots, n$ reveals aspects of the relationship not easily discerned by analyzing formulas or perusing numerical values. An example mentioned by Philip J. Davis[2] points this out:

> Students fooling around with . . . orbits in the many-body problem that have been graphically displayed have found periodic solutions whose existence defies our keenest analytical analysis.

In a very basic way, exploiting this extremely effective means of representing and conveying information about relationships is what this issue of *Computer* is all about. A graph is a visual image created using *graphics* (points, lines, etc.) and *geometry* (relative positions of the points, lines, etc). Combining these two well-developed fields can yield some powerful tools to aid scientists and engineers in their efforts to better understand the physical world.

The graphs required by today's scientists and engineers are more complex than a mere curve in an *x-y* Cartesian coordinate system. The relationships of interest often involve multidimensional quantities in both the domain and range, and "graphing" such relationships presents quite a challenge. Nonetheless, the basic principles of graphics and geometry used to create simpler graphs also yield useful scientific visualization techniques.

The articles in this issue discuss many of the techniques in use today and many others being researched and developed. Some techniques seem natural and obvious because they employ subtle metaphors. For example, temperature distribution in a 3D volume is really invisible, but a color-coded volume rendering with blues for low temperatures and reds for high temperatures can convey a great deal to the combustion scientist. Other techniques use interrogation operators to detect interesting and important regions. This permits focusing of rendering efforts and the viewer's attention.

Advances in scientific computation allow increasing complexity of mathematical models and simulations. This results in a closer approximation to reality, which enhances the possibility of acquiring new knowledge and understanding. Tremendously large collections of numerical values being produced and collected contain a great deal of information. The problem is to convey all of this information to the scientist to effectively use human creative and analytic capabilities. This requires a method of communication with a high bandwidth and an effective interface.

We use computer-generated images and human vision in scientific visualization. Images with color, intensity, transparency, texture, and a myriad of other techniques can, if properly prepared and properly explained, convey a tremendous amount of information in a short period of time.

While images provide a highly leveraged means of communication, sequences of images can yield more than the sum of their parts. Animation is an important additional dimension of scientific visualization. The mind can glean information from animation that is virtually impossible to obtain by separately viewing still images.

Because conventional publication media cannot effectively demonstrate the impact of animation, *Computer* Editor-in-Chief Bruce Shriver suggested the compilation of a video companion to this special issue. (See box at right.)

A video illustrates preprogrammed animation well, but it lacks the ability to convey the impact of interaction, which is another important dimension of scientific visualization. Allowing the researcher to cooperate with the computer by directing the computations increases the chances for inspiration, insight, and understanding.

The potential benefits of graphics and human/computer interaction were recognized more than a quarter of a century ago by Steve Coons and his student Ivan Suthterland. In a 1964 video interview[3] Coons described Sketchpad (Sutherland's thesis topic):

> You will see a designer, effectively solving a problem step-by-step, and he will not,

0018-9162/89/0800-0010$01.00 ©1989 IEEE

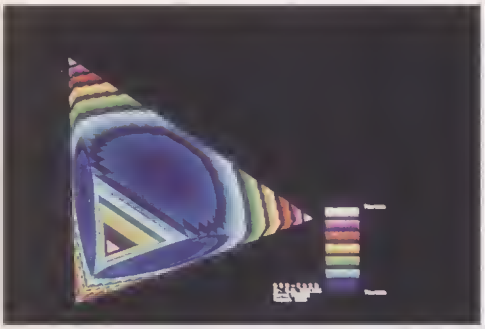

Far left: Flame data in vanishing cube. Left: Bezier volume. Images courtesy of Bernd Hamann, Computer-Aided Geometric Design Group, Arizona State University.

Scientific Computing

Arizona State University

> at the outset, know precisely what his problem is, nor will he know exactly how to solve it. But little by little, he will begin to investigate ideas, and the computer and he will be in cooperation, in the fullest cooperation, in this work.

By drawing from many areas, some very old (such as geometry) and some relatively new (such as computer graphics), researchers are developing new and exciting tools that will allow us to more fully cooperate with the computer. Thus we may gain a better understanding of the world we live in.

The goal of this special issue is to inform you of these new research activities and the impact they are having on our approach to science and engineering today and will have in the future. I hope you enjoy reading the articles and viewing the accompanying images and video. ☐

Acknowledgments

I wish to thank the members of the Executive Committee of the IEEE Technical Committee on Computer Graphics for their help and support. Larry Rosenblum, who chaired the committee, deserves special acknowledgment for all of the work he put forth.

My own efforts toward this special issue were supported by the Department of Energy under contract DE-FG-02-87ER25041 to Arizona State University. I also received a great deal of help from the College of Engineering and Applied Sciences at ASU, especially Engineering Computing Services, directed by Bill Lewis. Early in the planning stages of this special issue, Mike Wozny was very helpful, and I appreciate this. During the process of putting this special issue together, it has been a real pleasure to work with Bruce Shriver and, better yet, to learn from him.

Special thanks go to all the authors and referees.

References

1. R.W. Hamming, *Numerical Methods for Scientists and Engineers,* McGraw-Hill, New York, 1962.
2. P.J. Davis, "Visual Geometry, Computer Graphics and Theorems of Perceived Type," presented at the Missoula Conf. on the Influence of Computing on Mathematical Research and Education, August, 1973.
3. S. Coons, video interview at MIT, 1964, in *Visualization: State of the Art*, ed., L. Herr, SIGGraph Video Review, Special Issue #30, ACM, 1988. Order no. 915-130 for VHS, 915-030 for 3/4-inch tape. *Visualization: State of the Art Update* is Special Issue #35.

Gregory M. Nielson is a professor of computer science at Arizona State University, where he teaches and does research in the areas of computer graphics, computer-aided geometric design, and scientific visualization. He has published and lectured widely on the topics of curve and surface representation and design, interactive computer graphics, and scattered data interpolation.

Nielson is on the editorial board of ACM's *Transactions on Graphics*, the *Rocky Mountain Journal of Mathematics*, and *Computer Aided Geometric Design*. He is a member of the Executive Committee of the IEEE Computer Society Technical Committee on Computer Graphics.

Nielson received his PhD from the University of Utah in 1970.

Nielson's address is Computer Science Department, Arizona State University, Tempe, AZ 85287-5406.

Video animates scientific visualization

A color VHS videotape from the IEEE Computer Society Press illustrates scientific visualization, including preprogrammed animation. The video depicts the material covered in this special issue, presented by the authors. Additional presentations cover material from another collection of articles that join the 11 here in a companion volume to be published in Fall 1989. **See ad p. 146 for details.**

Visualization

Expanding Scientific and Engineering Research Opportunities

Thomas A. DeFanti and Maxine D. Brown, University of Illinois at Chicago
Bruce H. McCormick, Texas A&M University

Computational science and engineering (CS&E) describes a researcher's use of computers to simulate physical processes. CS&E parallels the development of the two other modes of science: theoretical and experimental/observational.

In addition to new methodologies, new technologies or mathematical tools have spurred the scientific revolutions. For example, calculus allowed Newton to codify the laws of nature mathematically and develop analytic methods for solving simple cases. Similarly, the development of the von Neumann computer architecture gave scientists the ability to solve the discretized laws of nature for general and complex cases.

CS&E now relies heavily on scientific visualization to represent these solutions, enabling scientists to turn mountains of numbers into movies and graphically display measurements of physical variables in space and time. This article explores the convergence of science and visualization, in support of its successful growth and development.

Visualization holds great promise for computational science and engineering, provided we can meet the immediate and long-term needs of both toolmakers and tool users.

What is scientific visualization?

Computer graphics and image processing are technologies. *Visualization*, a term used in the industry since the 1987 publication of the National Science Foundation report *Visualization in Scientific Computing*,[1] represents much more than that. Visualization is a form of communication that transcends application and technological boundaries.

A tool for discovery and understanding. The deluge of data generated by supercomputers and other high-volume data sources (such as medical imaging systems and satellites) makes it impossible for users to quantitatively examine more than a tiny fraction of a given solution. That is, it is impossible to investigate the qualitative global nature of numerical solutions.

With the advent of raster graphics, researchers can convert entire fields of variables (representing density, pressure, velocity, entropy, and so on) to color images. The information conveyed to the researcher undergoes a qualitative change because it brings the eye-brain system, with its great pattern-recognition capabilities, into play in a way that is impossible with purely numeric data.

For example, an observer instantly sees the vortices, shock systems, and flow patterns in a visualization of a hydrodynamic calculation, while these same patterns are invisible in mere listings of several hundred thousand numbers, each representing field quantities at one moment in time. When computing a space-time solution to the laws of physics, the particular numeric quantities at each event in time-space are not important; rather, what is important is understanding the global structure of the field variables that constitute the solution and the causal interconnections of the various components of that solution.

A tool for communication and teaching. Much of modern science can no longer be communicated in print. DNA sequences, molecular models, medical imaging scans, brain maps, simulated flights

0018-9162/89/0800-0012$01.00 ©1989 IEEE

through a terrain, simulations of fluid flow, and so on, all need to be expressed and taught visually over time. To understand, discover, or communicate phenomena, scientists want to compute the phenomena over time, create a series of images that illustrate the interrelationships of various parameters at specific time periods, download these images to local workstations for analysis, and record and play back one or more seconds of the animation.

According to the visualization report, "We speak (and hear) — and for 5000 years have preserved our words. But, we cannot share vision. To this oversight of evolution we owe the retardation of visual communication compared to language. Visualization by shared communication would be much easier if each of us had a CRT in the forehead."[1]

Our CRTs, although not implanted in our foreheads, are connected to computers that are nothing more than extensions of our brains. These computers, however, might not be in the same room with us. They could be down the hall, across town, or across the country. Hence, the ability to communicate visually — and remotely — with computers and each other depends on the accessibility, affordability, and performance of computers and computer networks.

The visualization report recommends the development of a federally funded initiative providing immediate and long-term funding of both research and technology developments (see Table 1).[1] Research developments are the responsibility of tool users — experts from engineering and the discipline sciences who depend on computations for their research. Technology developments are handled by toolmakers — the visualization researchers who can develop the necessary hardware, software, and systems.

Tool users' short-term needs

Every researcher requires a personal computer or workstation on his or her desk connected with a remote supercomputer. However, not all scientists require the same level of computing power. Hence, a three-tiered model environment is beginning to emerge that categorizes visualization systems by such factors as power, cost, and software support.

Table 1. Recommendations for a national initiative on visualization in scientific computing.

	Short-term Needs	Long-term Needs
Tool users: Computational scientists and engineers	Funding to incorporate visualization in current research	Funding to use model visualization environments
Toolmakers: Visualization scientists and engineers	No funding necessary	Funding to develop model visualization environments

Table 2. Visualization facility three-tiered hierarchy.

	Model A	Model B	Model C
Hardware	Supercomputer or super image computer	Minisupercomputer or image computer	Advanced workstations (mini-/ micro- image computer
Bandwidth (potential interactive rates, bits/second)	$>10^9$	10^7-10^8	10^3-10^6
Location (where users interact with the display screen)	Machine room (at the center)	Laboratory on a high-speed local area network	Laboratory on a national/ regional network
Software (in addition to discipline-specific data generation and processing)	Commercial packages for output only (no steering). Research required to develop interactive steering capabilities	Commercial packages are mostly output only. Some interaction is becoming available. Research required to improve discipline-specific interaction	Commercial packages and tools are widely available for both computation and interaction. Research required in languages, operating systems, and networking
Administration			
Strength:	Support staff	Discipline-specific visualization goals	Decentralization
Weakness:	Centralization	Small support staff	No support staff

Workstations. Researchers need workstations with access to supercomputers for

- immediate access to local graphics capabilities,
- networked access to supercomputers, and
- hard-copy recording.

Local graphics. Workstations, minicomputers, and image computers are significantly more affordable than supercomputers, and they are more powerful and effective visualization tools. There are already some 20 million personal computers and workstations in the United States, compared with about 200 supercomputers. Workstation users are increasingly treating supercomputers as one of many windows on the screen, and scientists must be able to "cut and paste" between the supercomputer and applications running on their local machines.

Access to supercomputers. Scientists need to transfer data to and from a main computation device, but today's networks are too slow for use in visualization. Some temporary techniques reduce the demand for high bandwidth, such as off-peak image transmission, image compression, image reconstruction from abstract representations, and local image generation. Networking is therefore as critical as computer power in helping scientists.

Hard-copy recording. Whether the visuals are for personal analysis, information sharing among peers, or presentations in formal surroundings, equipment for producing photographs, slides, videotapes, or laser disks needs to be in place and as easy to use as sending text files to a laser printer.

Scientists need the ability to create ad hoc graphics to verify the integrity of their simulations, gain insights from their analyses, and communicate their findings to others. Low-cost animation facilities should be connected to every user workstation so researchers can make scientific "home movies" with little effort. High-end visualization capabilities and facilities also should be available at all research centers; high-end graphics become important for presentation and publication of

Low-cost, visualization-compatible workstations and networks

Our ability to communicate visually and remotely with supercomputers and each other depends on

(1) the ease with which we can use our office/home computers to connect with the outside world, receive and transmit visual information, and record this information on videotapes or slides, and
(2) the cost/performance of today's networks.

The Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago is doing research in both areas. We are designing as our scientific animation workstation a low-cost computer system with a well-integrated visualization programming environment.

Users at the National Center for Supercomputing Applications (NCSA) — or any of the National Science Foundation-funded supercomputer centers, for that matter — cannot do graphics remotely due to slow network speeds, centralized and expensive graphics equipment, lack of graphics software tools, and the need for specialists in film/video production. Our research is motivated by the recent availability of low-cost graphics hardware and a good PC-based visualization toolkit, coupled with a growing awareness that scientists need visualization more for personal/peer analysis than for presentations.

EVL is integrating affordable commercial equipment with specially designed graphics software to make visualization a reality for computational scientists — whether they use their computers on a stand-alone basis or connected to supercomputers over networks. (Regarding affordability, academicians can generally receive $10,000 in equipment monies from their departments or colleges without applying for external grants — our yardstick is that equipment should cost no more than a three-year-old Buick.)

EVL's scientific animation workstation, shown in the accompanying figure, has hard-copy recording capability and an easy-to-use visualization environment to facilitate scientists' needs. The following list corresponds to items 1-6 in the figure.

(1) Supercomputer access. Supercomputers are most efficiently used to run complex simulation codes, the output of which is numbers. With access to graphics, researchers can convert numbers to pictures to qualitatively examine the global nature of their simulation output. Graphics can be made available on the host machine or, more efficiently, on the local workstation.

(2) Televisualization: graphical networking. As images require more colors, higher resolution, or larger volumes of data, they need more memory and become more impractical to transmit over networks or phone lines, to store on disks, or to convert and display on different frame buffers. EVL's Imcomp compression and conversion software converts images consisting of 24, 16, or 8 bits per pixel to 16 or 8 bits per pixel, then compresses them further to 2 or 3 bits per pixel while maintaining a reasonable full-color representation.[1,2] The program takes only 0.4 seconds to run on the Cray X-MP at NCSA, and it converts and transmits a 512 × 512 × 24-bit image from NCSA to EVL over a 56-kilobyte line within a few seconds.

Moreover, visuals must be transmitted from memory to memory (that is, from supercomputer memory to frame buffer memory in the local computer), not just from file to file as in electronic mail-type networks. NCSA's Telnet communications software has been modified to do this and expanded to include Imcomp routines that automatically compress images.

In addition to compression, value-added nodes speed up graphical transmission by balancing transmission costs with local computing costs. Model data is sent over networks and then rendered or reconstructed at the scientist's end. EVL is currently investigating the use of its AT&T Pixel Machine as a graphics server that would render model data transmitted over the network from the supercomputer and then transmit the resulting images over a local area network to individuals' desktop computers.

(3) Truevision Vista graphics board. Scientists need to be able to preview, record, and play back animations at any speed and in cyclical fashion to examine the dynamics of their data changing over time, to spot anomalies, or to uncover computation errors. The Vista board's large configurable memory allows us to get anywhere from 32 screens at 512 × 512 pixels to 128 results once researchers conclude their work.

Three-tiered model computational environment. Observations of the way scientists use visualization suggest that a three-tiered model environment is evolving, as defined in Table 2. Each model is distinguished by hardware costs, computing power, bandwidth, location, software support, and administrative considerations.[1]

This model environment assumes that scientists want as direct a visual connection to their computations as possible. While supercomputers (model A) provide scientists with powerful number-crunching tools for generating data, they currently do not produce graphics; they do fill arrays with information that somehow gets piped to display devices. (Table 2 assumes that supercomputers and super image computers have equivalent power. Super image computers, although not commercially available today except in the form of a special-purpose flight simulator, will provide the specialized processing necessary for real-time volume visualization.)

Workstations give scientists more control over their visual output (models B and C). A workstation typically addresses its display memory the same way it addresses regular memory, incurring essentially no hardware overhead to display computed results. (Table 2 also assumes that minisupercomputers and image computers have equivalent power, and that advanced workstations and mini-/micro- image computers have equivalent power.)

Scientists should be able to select either more-expensive workstations with powerful visualization potential (model B) or less expensive ones (model C) while maintaining network connections to larger machines (model A) to do computations when necessary. This interdependency can work quite well. For example, a scientist can calculate 20-60 frames of a simulation sequence on a supercomputer, download the images to a workstation to create a minimovie, and then play back the sequence at any speed under local control.[2,3] (See sidebar, "Low-cost, visualization-compatible workstations and networks.")

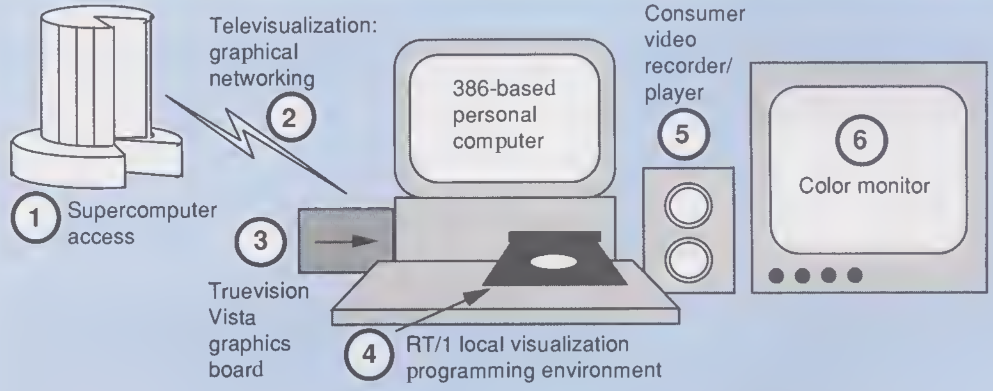
Televisualization: graphical networking
2
Consumer video recorder/ player
386-based personal computer
5
6
Color monitor
1
Supercomputer access
3
Truevision Vista graphics board
4
RT/1 local visualization programming environment

The Electronic Visualization Laboratory's RT/1 graphics language, an 80386-based personal computer, the Truevision Vista board, and consumer video gear comprise a scientific animation production facility that is economical enough to be made available to research scientists and engineers on a broad scale.

screens at 128 × 128 pixels, all at 8 bits per pixel. The board is also video compatible, so images can be recorded directly to videotape.

(4) Real Time/One (RT/1) local visualization programming environment. Scientists need a set of tools for picture composition, picture saving/restoring, fonts and text, resizing, rotation, moving, copying, hand retouching (painting), color manipulation, etc. They also need a local graphics programming environment in which to develop new tools or extend the capabilities of existing ones.

RT/1, an easy-to-use graphics programming language developed by EVL faculty and students, meets the criteria required of a visualization system environment. The language, written in C and running under Unix and MS-DOS, runs on all of EVL's workstations and personal computers. EVL is porting RT/1 to new workstations as they are acquired, extending the capabilities of the language, and developing application programs tailored to the needs of scientists.

(5) Consumer video recorder/player. If it's not recordable, it's not science. Moreover, the equipment for producing videotapes needs to be as easy to use as sending text files to a laser printer. We are integrating low-cost consumer video equipment into the workstation so scientists can quickly and easily preview and record frames of animation.[3] This equipment also comes with a built-in microphone so scientists can add narration or other sounds to visual recordings.

(6) Color monitor. Today's consumer video systems not only record but also can be attached to any television for immediate viewing of recorded material. Scientists can take a small video unit to a conference and plug it into a television there to share findings with colleagues. Should peers in other towns have similar equipment, colleagues could mail tapes to each other for viewing.

References

1. M.D. Brown and M. Krogh, "Imcomp — An Image Compression and Conversion Algorithm for the Efficient Transmission, Storage, and Display of Color Images," *NCSA Data Link*, Vol. 2, No. 3, National Center for Supercomputing Applications, June 1988, pp. 11-24.
2. G. Campbell et al., "Two-Bit/Pixel Full Color Encoding," *Computer Graphics* (SIGGraph Proc.), Vol. 20, No. 4, Aug. 1986, pp. 215-223.
3. T.A. DeFanti and D.J. Sandin, "The Usable Intersection of PC Graphics and NTSC Video Recording," *IEEE Computer Graphics and Applications*, Oct. 1987, pp. 50-58.

Table 3. Total corporate computing needs. (Source: Larry Smarr, NCSA, Sept. 1988.)

	Computational Science and Engineering	Data Processing
Corporate officer responsible	Vice president of research or long-range planning	Vice president of management information systems (MIS)
Tiered architectures	Personal computers and graphics workstations; midrange machines (mainframes/minisupercomputers; supercomputers *Need exists for multivendor, networked, hierarchical computing.*	Personal computers; minicomputers; mainframes *Software portability only partially exists between these levels, and then only within one vendor's product line.*
	Open systems *No vendor has emerged who offers integrated systems and end-to-end solutions. As a result, end users are faced with a confusing set of products from various vendors and nowhere to turn for advice on how to integrate them.*	Closed systems *IBM and Digital Equipment Corporation manufacture all levels of computers and the connections between them.*
Vendors	Fragmented market populated by start-ups and extremely high-growth companies: Workstations (Sun, DEC, Apollo, IBM, Hewlett-Packard, Apple, Silicon Graphics, Ardent, Stellar, AT&T Pixel, etc.); Midrange (DEC, IBM, Alliant, Amdahl, Convex, Scientific Computing Systems, Multiflow, Elxsi); Supercomputers (Cray, IBM).	Mature, slow-growth marketplace dominated by a few giant vendors, such as IBM and DEC.
Operating Systems	Unix	MVS, DOS, VMS (proprietary)
Networking protocols; telecommunications; speeds	Open network standards; Long-haul telecommunications; High speed = 1,000 Mbits/second *Because of the scarcity of $20 million supercomputers, most universities and corporate CS&E users are remote and must gain access to supercomputers over long-haul telecommunication lines.*	SNA, DECnet (proprietary); High speed = 50 Mbits/second *Within a corporation, most networks hook many "dumb terminals" up to a central mainframe where all the computing power resides. PCs are generally used stand-alone; those networked to a mainframe generally use the network to download or upload files, and computing is decoupled.*
Common unit of information	Image (megabyte) *Supercomputer simulations produce such enormous amounts of data that visualization is essential.*	Number (byte)
Common unit for computation speed	Mflops	MIPS

Additional models D, E, and F, corresponding to personal computers, alphanumeric CRT terminals, and batch output, respectively, also exist. They do not represent advanced visualization technology, so they are not included in our model environment. Note, however, that model F has been used to produce a great deal of animation for both the scientific and commercial entertainment industries for the past 20 years.

Tool users' long-term needs

CS&E is emerging as a new marketplace with needs distinct from those of data processing, as shown in Table 3. Success in the CS&E marketplace of the 1990s will depend on a commitment to standards, ease of use, connectivity, open systems, integrated systems, software portability, multivendor environments, leading-edge technology, and customer service and support.[4]

The list of research opportunities for visualization in scientific computing is long and spans all of contemporary scientific endeavor. The sidebar "Scientific and engineering research opportunities" presents specific examples of advanced scientific and engineering applications to show

(Continued on p. 22)

Scientific and engineering research opportunities

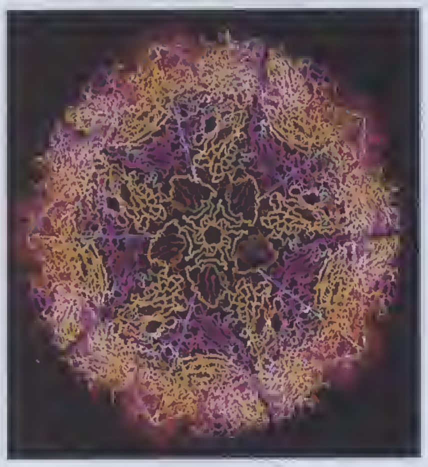

Molecular modeling. The use of interactive computer graphics to gain insight into chemical complexity began in 1964. Interactive graphics is now an integral part of academic and industrial research on molecular structures, and the methodology is being successfully combined with supercomputers to model complex systems. Two types of images can currently be generated: realistic pictures of molecules and 3D line drawings. Raster equipment is used to create realistic representations and animations, while vector hardware, used for real-time display and interaction, creates line drawings.

The image at left is a 3D line drawing of the rhinovirus, the common cold virus, showing its geometric structure and complexity. The image at right is an artistic rendering of the human papilloma virus (HPV). It was done by a group of Chicago-area artists who appreciate the underlying mathematics of nature and the complexity of the inner workings between atoms.[1]

Left-hand © 1988 T.J. O'Donnell. Data courtesy of Dr. Rossman, Crystallography Group, Purdue Univ. Image courtesy of the EVL, Univ. of Illinois at Chicago. Right-hand © 1989 (Art)^n Laboratory, Illinois Institute of Technology. (Art)^n artists: Donna Cox, NCSA, Univ. of Illinois at Urbana-Champaign; Stephan Meyers, Dan Sandin, and Tom DeFanti, EVL, Univ. of Illinois at Chicago; Ellen Sandor, (Art)^n Laboratory, Illinois Institute of Technology.

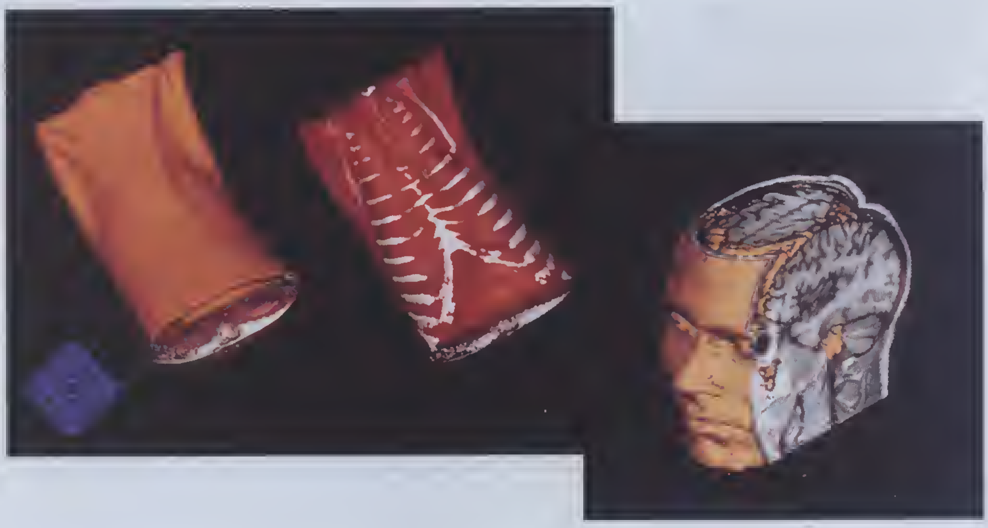

Medical imaging. Scientific computation applied to medical imaging has created opportunities in diagnostic medicine, surgical planning for orthopedic prostheses, and radiation treatment planning. In each case, these opportunities have been brought about by 2D and 3D visualizations of portions of the body previously inaccessible to view.

The above-left image is a shaded surface volume rendering of a 128 × 128 × 197 computerized tomography scan of a tree sloth. The opacity of various structures can be interactively modified to show the skin surface or to reveal internal structures. The bones of the rib cage, shoulder blades, and spine can be seen in the image on the right, as well as the trachea, lungs, heart and diaphragm.

The above-right image is a shaded surface volume rendering of a 256 × 256 × 61 magnetic-resonance imagery (MRI) scan of a human head. The rendering shows a mixture of surface and slice-based techniques, where external structures such as the skin are rendered with surface shading, while slice planes are voxel-mapped to reveal the original MRI data. Physicians can use this technique to relate the position of internal structures such as tumor sites to external landmarks. These images were generated using the Voxvu volume rendering tool on a Sun workstation with the TAAC-1 Application Accelerator.

© 1989 Chuck Mosher and Ruth Johnson, Sun Microsystems. Data for above-left image courtesy of Eric Hoffman, UPA. Data for above-right image courtesy of Jeff Shaw, Vanderbilt Univ.

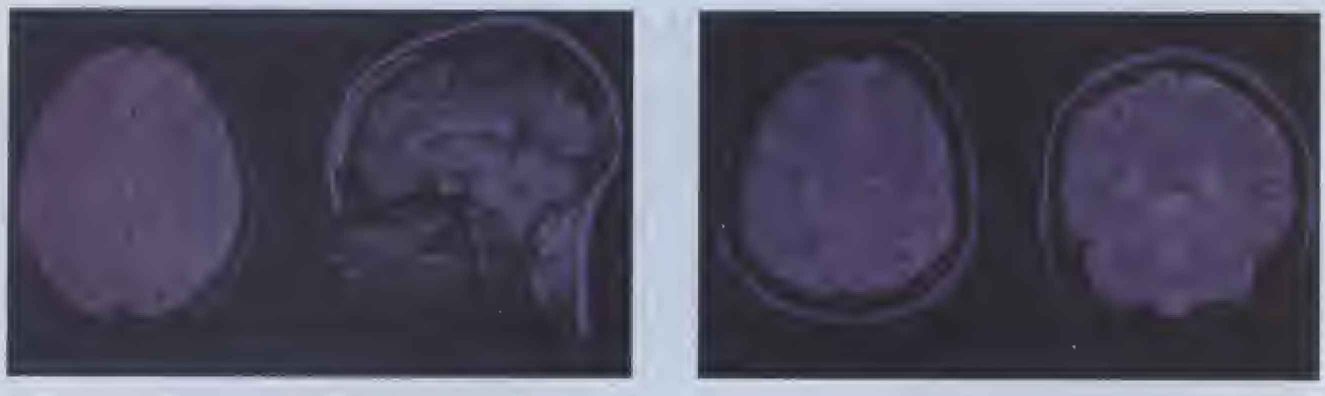

Brain structure and function. Rutgers University is using computer vision and visualization methods to automatically detect white-matter lesions in MRI scans of the human brain.

In the above-left image, low-level vision methods locate the outline of the brain, landmarks such as the interhemispherical fissure plane, and suspected lesions. The system calculates the orientation of the brain and uses the segmentations provided by the low-level methods to fit a deformable model to each patient's brain to determine the position and shape of difficult-to-identify organs or regions of interest. This customized model, shown in the above-right image, is used to obtain information about the anatomical position of the suspected lesions so that the system can reject false positives and determine the affected organs. The system has been tested on more than 1,200 images from 19 patients, producing good results.

© 1989 Ioannis Kapouleas, Computer Science Dept., Rutgers Univ.

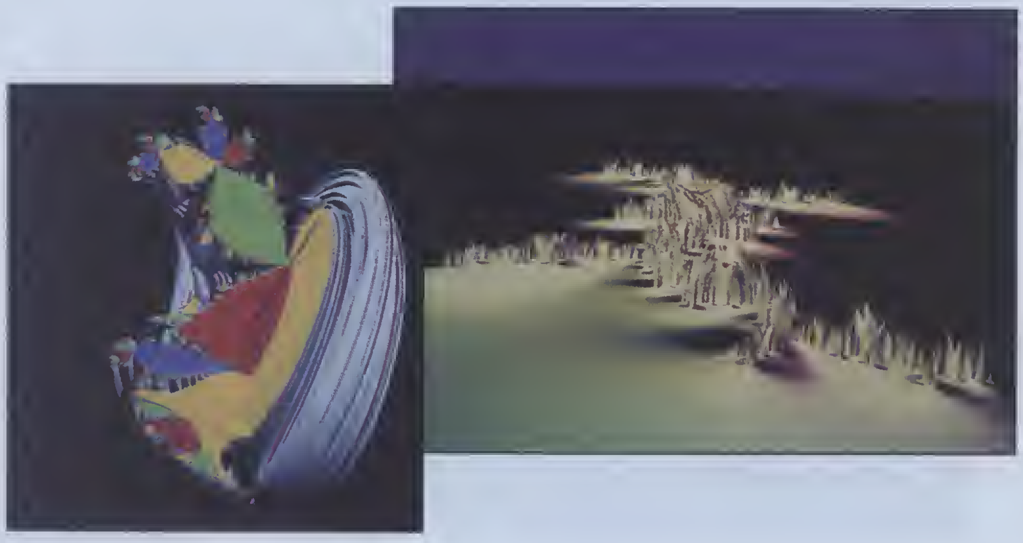

Mathematics. These images illustrate a type of fractal known as the Julia set. A filled-in Julia set is a set of points that do not converge (or diverge) to infinity after repeated applications of a function, such as $f(z) = z^2 + c$. These functions are often investigated in the complex plane, but they also exist in the quaternions, a coordinate system that spans one real and three imaginary axes. Visualization helps mathematicians understand these equations, which are too complex to conceptualize otherwise.[2]

The above-left image is a quaternion filled-in Julia set minus its front-upper-left octant; the inner components of the four-cycle are revealed, defining its basin of attraction. The above-right image is a visualization of a dendritic quaternion Julia set in the complex plane; the unusual lighting uses a 3D gradient in the complex plane.

© 1989 John Hart, EVL, Univ. of Illinois at Chicago.

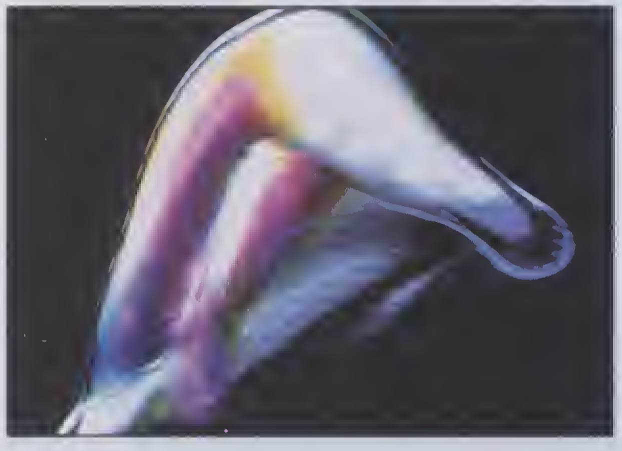

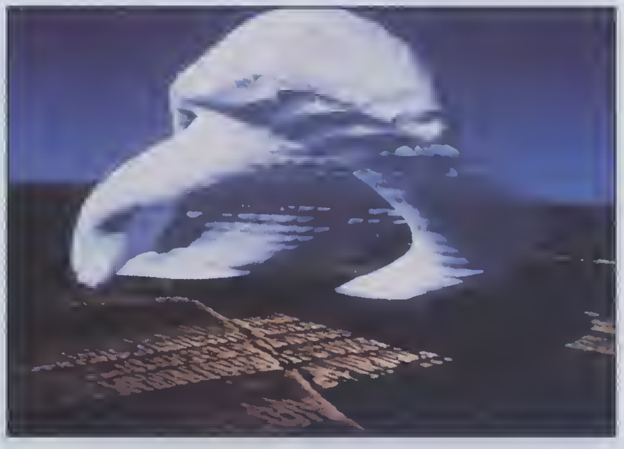

Geosciences: meteorology. The study of severe storms through observation and modeling helps research meteorologists understand the atmospheric conditions that breed large and violent tornadoes and the mechanisms by which tornadoes form and persist.[3] Theoreticians and field workers obtain information on behavior that cannot be safely observed; study the interactions of various environments, characterized by differing vertical wind, temperature, pressure, and moisture structures; and obtain useful guides for future research.

Transparency and volumetric rendering are used to view multiple surfaces; shading is used to display individual solid surfaces.

The above-left image uses voxel (grid cell) data to display rainwater and vertical vorticity information about a storm,[4] providing scientists with more information than if they had observed the storm with their eyes. The fuzzy region indicates low rainwater amounts while the bright white regions indicate large amounts of rainwater within the cloud. The vertical vorticity is texture mapped onto the rainwater with color; purple indicates dominant positive vorticity and blue indicates dominant negative vorticity.

The above-right image is from an animated simulation of a storm over Kansas, in which the rainwater surface was polygonized (tiled) and then rendered. The simulation clearly reveals substantial variations in the structure of the rainwater field not apparent earlier.

Above-left © 1988 Robert Wilhelmson and Craig Upson, NCSA, Univ. of Illinois at Urbana-Champaign. Above-right © 1988 Robert Wilhelmson, Crystal Shaw, Lou Wicker, Stefen Fangmeier, and the NCSA Visualization Production Team, Univ. of Illinois at Urbana-Champaign.

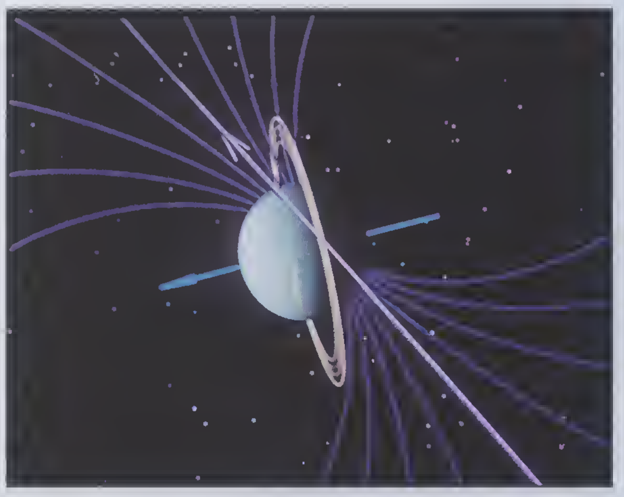

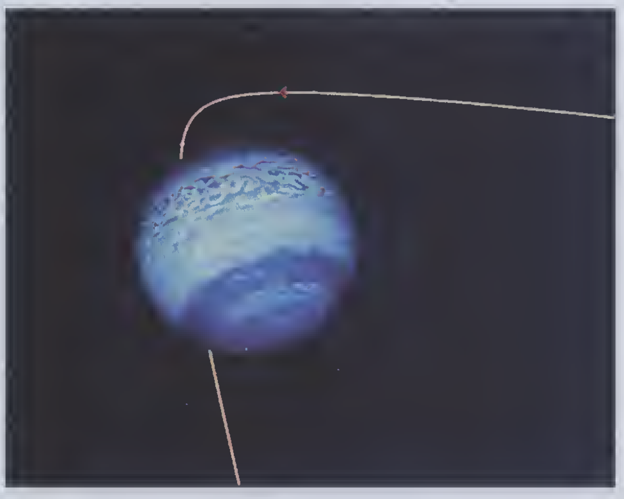

Space exploration. The field of planetary study involves the accumulation of huge volumes of data on the planets in the solar system. Enough data is now available that scientists are beginning to integrate observed phenomena and theory from other fields involved in planetary study: meteorology, geography, planetary physics, astronomy, and astrophysics.

The above-left image is from an animated simulation of the dynamics of Uranus' magnetosphere. The simulation shows that the angle of the dipole axis (purple arrow) is offset from the planet's angle of rotation (aqua arrow). The above-right image is from a simulation of the Voyager 2 Neptune encounter to occur in late summer of 1989. This image illustrates the path of the Voyager 2 as viewed from Earth.

Above-left © 1989 Computer Graphics Group of the Jet Propulsion Laboratory and G. Hannes Voigt of Rice Univ. Above-right © 1989 Computer Graphics Group of the Jet Propulsion Laboratory.

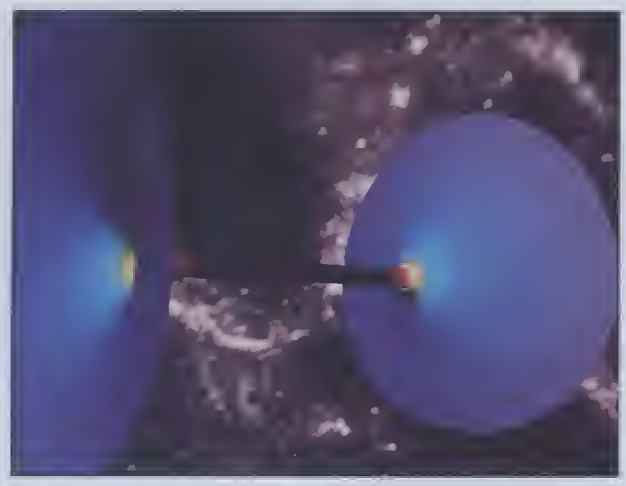

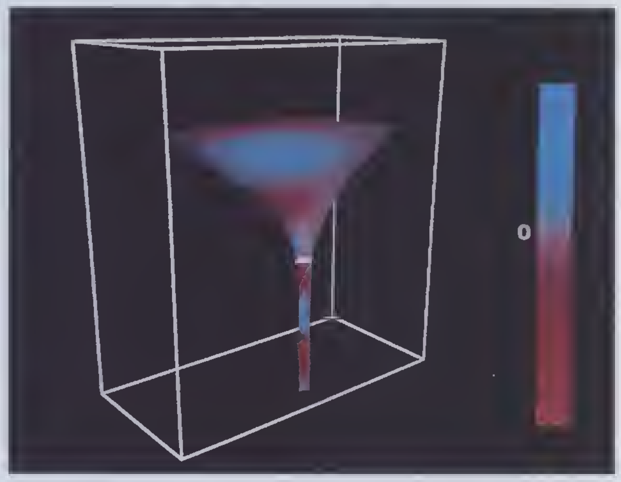

Astrophysics. Computational astrophysicists at the NCSA work with artists in an attempt to see the unseen and create visual paradigms for phenomena that have no known visual representation.

An embedding diagram of a Schwarzschild black hole and the behavior of its gravitational field, illustrated in the above-left image, was obtained from a numerical solution of Einstein's numerical relativity equations. The surface of the diagram measures the curvature of space due to the presence of the black hole, while the color scale represents the speed at which idealized clocks measure time (with red representing the slowest clocks and blue representing the fastest).

A black hole emits gravitational radiation after it has been struck by an incoming gravity wave. The above-right image is from an animated sequence that shows, for the first time, the influence of the curved space on the propagation of the radiation. Through the use of an isometric embedding diagram, the curvature of the space surrounding the black hole is represented by the surface on which the waves propagate. The white ring locates the surface of the black hole, and the regions above and below represent the exterior and interior of the black hole, respectively.

© 1989 David Hobill, Larry Smarr, David Bernstein, Donna Cox, and Ray Idaszak, NCSA, Univ. of Illinois at Urbana-Champaign.

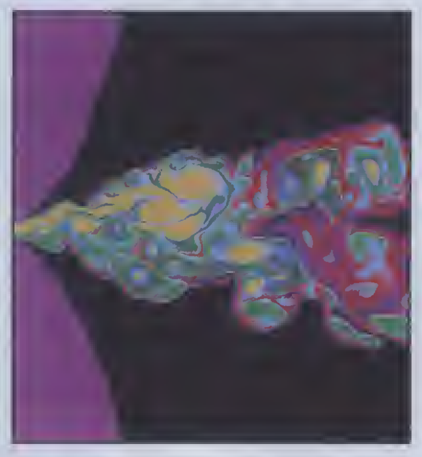

Computational fluid dynamics. Computational astronomers rely on supercomputing and visualization techniques to understand why jets from some galaxies flare dramatically. Magnetohydrodynamics code is used to solve equations that describe the flow of a fluid or gas with magnetic fields using finite differences.

The above image is a visualization of a cosmic jet traveling at Mach 2.5 passing through a shock wave (located at the left of the image). The jet abruptly slows and breaks up into a broadened subsonic plume whose morphology, or shape, is strikingly similar to that of a radio lobe of a wide-angle tailed galaxy. The morphology of the jet after impact is emphasized through the use of pseudocolor. This research has given astronomers important clues about why jets from some radio galaxies flare into broad plumes while jets from others remain remarkably straight and narrow.[5,6]

© 1989 Michael Norman and Donna Cox of the NCSA, Univ. of Illinois at Urbana-Champaign, and Jack Burns and Martin Sulkanen of the Univ. of New Mexico.

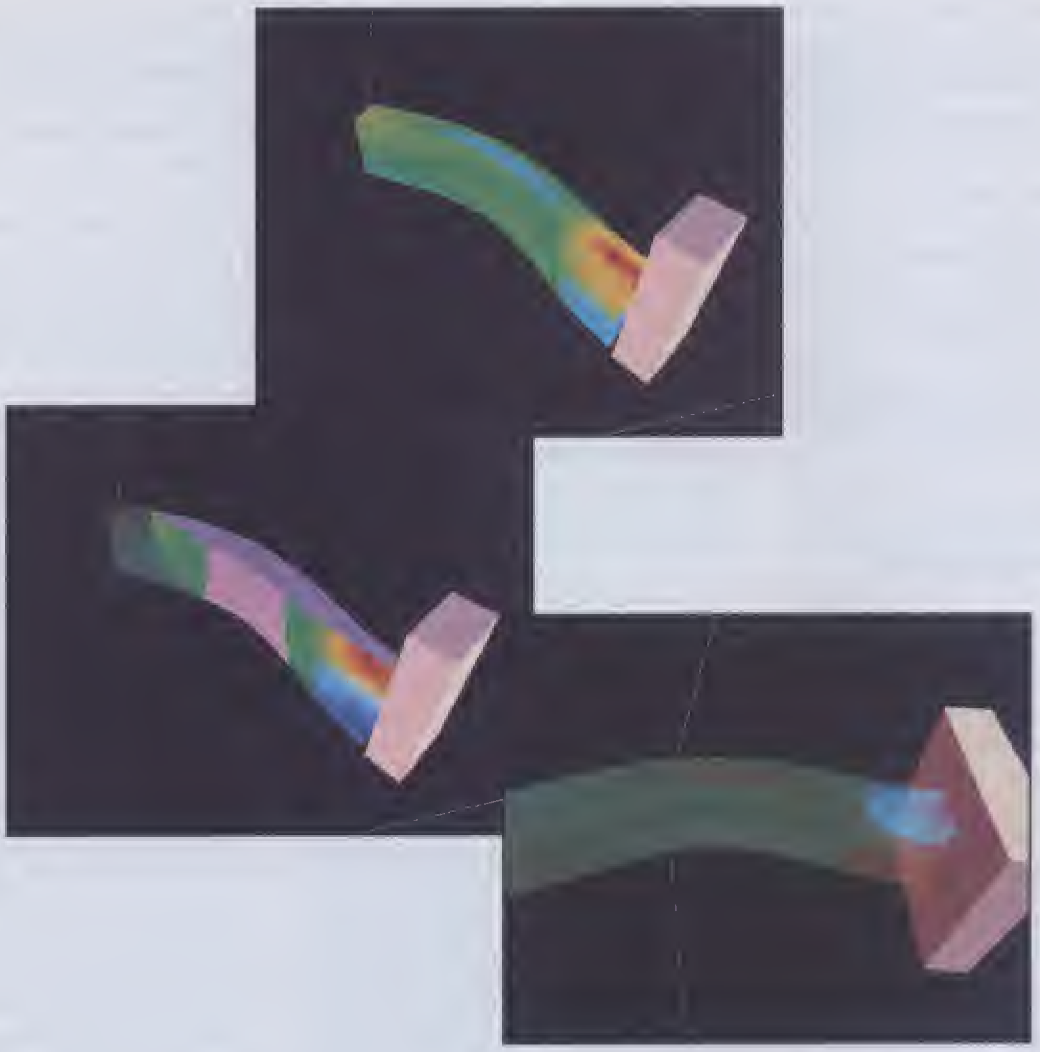

Finite element analysis. Finite element analysis is used in this example to show the stress distribution in a beam at its maximum tip displacement in the third eigenmode. The results were computed using linear elastic elements and a lumped-mass approximation.

The top image uses a conventional approach of displaying the stress values on the outer surface of the deformed shape. The middle image uses a cutting plane to look at the stress values on a cross section of the root of the beam. The bottom image shows a different view of the beam and uses an iso-contour stress surface to convey the three-dimensional nature of the stress concentration at the root of the beam. These images are still frames from fully animated and interactive models. They were computed and rendered on a Silicon Graphics 4D/120 GTX workstation using the SolidView program to perform real-time cutting and iso-contour surface generation.

© 1989 James M. Winget, Silicon Graphics.

References

1. R.P. Feynman, *Surely You're Joking, Mr. Feynman! Adventures of a Curious Character*, Bantam Books, 1986, pp. 236-253.
2. J.C. Hart, D.J. Sandin, and L.H. Kauffman, "Ray Tracing Deterministic 3D Fractals," to be published in *SIGGraph 89 Conf. Proc., Computer Graphics*, Vol. 23, No. 4, Aug. 1989.
3. R.B. Wilhelmson, "Numerical Simulations of Severe Storms," *Proc. Fourth Int'l Symp.: Science and Engineering on Cray Supercomputers*, Cray Research, Oct. 1988.
4. C. Upson and M. Keeler, "VBuffer: Visible Volume Rendering," *SIGGraph 88 Conf. Proc., Computer Graphics*, Vol. 22, No. 4, Aug. 1988, pp. 59-64.
5. K.-H.A. Winkler and M.L. Norman, "Munacolor: Understanding High-Resolution Gas Dynamical Simulations Through Color Graphics," *Astrophysical Radiation Hydrodynamics*, D. Reidel Publishing, 1986, pp. 223-243.
6. N.J. Zabusky, "Computational Synergetics," *Physics Today*, July 1984, reprint.

(Continued from p. 16)

how visualization tools are helping researchers understand and steer computations. Our examples fall into the following categories:

- Molecular modeling,
- Medical imaging,
- Brain structure and function,
- Mathematics,
- Geosciences (meteorology),
- Space exploration,
- Astrophysics,
- Computational fluid dynamics, and
- Finite element analysis.

Toolmakers' short-term needs

Commercial industry currently supports visualization hardware and software, as listed below. There is a pressing need to educate the scientific and engineering research communities about the available equipment.

Software. Commercial visualization software exists in the following categories:

Lines. The earliest software for graphics drew lines in three dimensions and projected them onto a two-dimensional plane, offered viewing transformations for looking at the result, and offered transformations (scale, rotate, and translate) for describing the line objects. The theory and practice of drawing lines, expressed in homogeneous coordinates, and the control and display of lines using 4 × 4 matrices, represented a major development in computer graphics.

A variety of current standards incorporate these basic principles, and the CAD/CAM industry has embraced this level of the art. It is cheap enough to put on every engineer's desk and fast enough for real-time interaction.

Polygonal surfaces. The next level of software — surfaces represented by polygons — has only recently been built into hardware. Polygon filling, or tiling, is commonly available in hardware and software. Hidden surface removal is included, and antialiasing of polygon edges is sometimes provided to remove distracting stairsteps, or jaggies. Light sources can be incorporated into the rendered image, but they are usually point sources at infinity emitting white light.

Patches. The next level of sophistication represents surfaces as curved surface pieces called patches. This is still largely a software domain, although we expect hardware to appear soon. The most advanced software packages handle a variety of patch types. They also provide very sophisticated lighting models with multiple-colored lights and distributed or point sources located either at infinity or in the scene.

Antialiasing is assumed, and the packages handle optical effects such as transparency, translucency, refraction, and reflection. Research software provides even more features that produce greater realism, such as articulated motion blur, depth of field, follow focus, constructive solid geometry, and radiosity.

The software contains no practical limit on scene complexity (such as the number of allowable polygons), but computation of highly complex scenes on a supercomputer can take anywhere from 0.5 to 1.5 hours per frame.

Image processing. Image processing software has followed a separate path over the last 15 years. The elaborate software packages now available provide functions such as convolution, Fourier transform, histogram, histogram equalization, edge detection, edge enhancement, noise reduction, thresholding, segmentation, bicubic and biquadratic warping, and resampling.

Many of these functions have been hardwired into special boards. General-purpose processors have only recently become powerful enough to make software competitive with hardware while maintaining generality. Image computers can run both computer graphics and image processing software packages.

Animation. In its broadest sense, animation means movement. It frequently connotes the complex motion of many objects, possibly articulated, moving simultaneously, and interacting with one another. Animation is desirable for the visualization of dynamic, complex processes. Basic animation control routines should be part of any standard visualization tool kit.

Glue. A class of software appreciated by visualization professionals but not necessarily by scientists is the "glue" used to combine images generated or analyzed by the packages described above. For convenience, a user must have tools for picture composition, picture saving/restoring, fonts and text, resizing, rotation, moving, copying, hand retouching (painting), color manipulation, etc. Together, these functions comprise a visualization environment system, which is to visualization what an operating system is to general computing.

Window systems. Windowing systems are commonplace in black-and-white bit graphics and are being extended to color graphics. Visualization software must incorporate and remain consistent with windowing concepts.

Volume visualization. Volume visualization software is still rudimentary. Algorithms for rendering lines, curves, surfaces, and volumes into volume memories are only now becoming available.[5,6] Hidden volume removal is unknown, the compositing of volumes is yet to be fully addressed, 3D paint programs (sculpting programs) have yet to be written, and general utilities for arbitrary rotation and size change of volumes do not exist. In other words, there is much research to be done in this field.

Hardware. The following are available commercial visualization hardware tools:

Input devices. Current digital input devices include supercomputers, satellites, medical scanners, seismic recorders, and digitizing cameras. The rapidly increasing bandwidth of these devices emphasizes the need for work in volume visualization.

We expect continued improvement in the resolution and bandwidth of input devices. Supercomputers will get faster and the resolution of images from satellites will increase. Real-time video digitizers already exist. Monochrome digital digitizers with 2,048 × 2,048-pixel resolution are becoming quite fast, although they do not yet operate at real-time speeds. Print-quality input scanners are still quite expensive, but we expect the prices to fall as digital technology cheapens and competing scanning technologies mature. CCD (charge-coupled device) array input scanners will improve in resolution and become serious candidates for input devices in high-resolution work.

Interactive input devices are continually improving. Common 2D devices include knobs, switches, pedals, mice, and tablets. Tablets are the most general and also need the most improvement; they need higher resolution, higher speed, and more degrees of freedom.

Six-dimensional interactive devices are also available, providing the usual 3D positional information plus three degrees of orientation information (yaw, pitch, and roll). Higher-dimensional devices, such as the data glove, have begun to appear. These will improve to offer higher resolution, higher speed, and lower cost.

Output devices. Raster displays of 2D frame buffers have improved steadily to offer more colors, higher resolutions, and less flicker. A typical color raster display today offers a 1,280 × 1,024-pixel display at 60 frames per second and 24 bits of color per pixel (16 megacolors).

High-definition television (HDTV) — a proposed standard that will offer larger, brighter, sharper pictures than currently available in video — will affect visualization. Also, video is moving toward an all-digital format, designated 4:2:2, to standardize digital interconnections of diverse video products.

Color raster displays will evolve toward 2,048 × 2,048 pixels in the next several years. The displays themselves already exist in limited quantities, but the computational bandwidth required to feed them is still lacking. Black-and-white 2D raster displays already have resolutions greater than 2,048 × 2,048 pixels with enough bandwidth to feed them. These displays will certainly reach even higher resolutions in the next five years.

Stereo displays are also beginning to appear commercially, and we confidently predict that these will improve in screen size, resolution, brightness, and availability. These displays will be quite helpful in volume visualization.

Other output devices are similarly improving. HDTV will spur the development of compatible recorders. Film recorders will become cheaper as the technology becomes cheaper and the competition matures. Should stereo become a widely accepted mode of presentation for volume visualizations, then stereo film and video standards will have to be developed.

Workstations. Fast vector machines are now common and have extensive use in such areas as CAD/CAM and real-time 3D design. Recently, they have improved to offer color vectors and perfect end-matching. Frame buffers have been added so that surface raster graphics can be combined with vector displays.

Also, fast surface machines are about to arrive. They exist in simplified forms already and in more advanced states as firmware in special machines. Chips are now being built to speed up certain aspects of surface rendering, particularly the tiling of polygons. By 1990, full hardware support of surface graphics will be available, offering rendering features such as texture mapping, bump mapping, antialiasing, reflections, transparency, and shadows.

Vector machines will initially serve as powerful, real-time front ends to surface machines. Eventually, surface machines will be cheap and fast enough to permit scientists to do real-time design using surfaces rather than lines.

Among image processors, fast planar machines have existed for some time. These machines contain special boards for certain aspects of image processing, such as fast Fourier transforms. Faster versions are becoming available that have wider processing capabilities and higher resolutions. In fact, the notion of a general-purpose image processor that can implement any image processing algorithm as a program is becoming common.

Toolmakers' long-term needs

Raw computing power would be more effectively harnessed than it is today if calculations could be understood pictorially and their progress guided dynamically. Modern modes of computing involve interactive, extemporaneous generation of views from masses of data and exploration of model spaces by interactive steering of computations.

A scientist's ability to comprehend the results of his or her computations depends on the effectiveness of available tools. To increase that effectiveness, we need to

- encourage the production of documented, maintained, upward-compatible software and hardware;
- motivate manufacturers to solve network bottleneck problems;
- encourage universities to incorporate CS&E and visualization in computer science, engineering, and discipline-science curricula; and
- guarantee the publication and dissemination of research and results on a variety of media.

Hardware, software, and systems. General visualization issues that need to be supported include:

- Interactive steering of simulations and calculations
- Workstation-driven use of supercomputers
- Graphics-oriented programming environments
- Higher-dimensional visualization of scalar, vector, and tensor fields
- Dynamic visualization of fields and flow
- High-bandwidth picture networks and protocols
- Massive data-set handling, notably for signal and image processing applications
- Vectorized and parallelized algorithms for graphics and image processing
- Specialized architectures for graphics and image processing
- A framework for international visualization hardware and software standards

Networking. The application of networks to visualization, called *televisualization*, encompasses much more than text transfer (such as electronic mail) and gateway protocol decoding. It also involves image transfer, which entails compression, decompression, rendering, recognizing, and interpreting. Televisualization requires a major enhancement over existing network capabilities in the following areas:

Increased data rates. The sheer scale of graphics and imaging data sets challenges the current bandwidth and interactivity of networks. Networks handle screenfuls of textual information well; network nodes are simply gateways that neither add nor detract from the quality of the message. But a 512×512-pixel image with 8 bits per pixel has approximately 100 times more information than a screen of text with 25 rows and 80 characters per row. A 1,024 × 1,024 × 1,024-voxel volume with 48 bits per voxel contains 16,000 times more information than a 512 × 512-pixel image. Gigabit speeds are sufficient to pass volumes of the current size of 256 × 256 × 256 voxels with 4 bytes per voxel, but this rate will have to be extended within several years to 1-gigabyte/second channels.

Compression/decompression algorithms. Compression improves the speed with which visual data is transmitted. Current schemes for full-color image compression work well,[7] but other forms of compression must be researched, and comprehensive protocols must be developed for managing all these capabilities:

Table 4. The evolution of communication tools.

Communications media	Number of years old
Sight	5×10^8
Speech	5×10^5
Writing	5×10^3
Print broadcasting	5×10^2
Visual broadcasting	5×10^1
Visualization	5×10^0

- Transmit the procedures to create the images rather than the images themselves.
- Transmit endpoints of vector images.
- Transmit polygonal, constructive solid geometry, or bicubic patches of surface models.
- Transmit semantic descriptions of the objects.

Value-added processing at nodes. Value-added nodes also speed up graphical transmission. Computers process text and numbers in main memory, occasionally transmitting some of them to peripherals. Images, however, often must be transferred to special memories for rendering, 3D imaging, or viewing. Each instance of transferring and processing an image aims to increase its visualization value to the scientist. The ability to process images at various nodes along a network embraces the central concept of distributed processing.

In distributed computing, transmission costs are balanced with local computing costs. It sometimes makes more sense to send model data over networks and then render or reconstruct the data at the scientist's end. This presumes that there is appropriate equipment at both ends, that the various software modules are compatible with one another, and that the software can run on a variety of equipment types.

A televisualization network for image passing between machines is analogous to the software paradigm of message passing between process layers. This type of networking, combined with interaction, cannot be achieved using conventional Fortran subroutine calls. Significant software development and protocol standardization are necessary to bring televisualization to the discipline sciences.

Interaction capabilities. Interactive visual computing is a process whereby scientists communicate with data by manipulating its visual representation during processing. The more sophisticated process of navigation allows scientists to dynamically modify, or steer, computations while they are occurring. This lets researchers change parameters, resolution, or representation, and then see the effects.

Teaching CS&E and visualization. The principal barrier to growth in the CS&E market is the fact that corporate researchers and managers lack education and training in CS&E technologies and methodologies. Few industrial researchers know how to use distributed CS&E to do their work and, more importantly, they do not know how to think computationally and visually. Other roadblocks include the following:

- The Association for Computing Machinery's approved computer science curriculum lists computer graphics as merely one of many optional topics; image processing is not mentioned at all.
- Engineering accreditors do not require computer graphics or image processing.
- Many engineering school deans are unaware of the importance of visualization or cannot justify the hardware and software expense involved in teaching the subject.
- The number of tenured faculty teaching computer graphics in American universities is about the same today as 15 years ago, and they are roughly the same people.
- Scientists, while educated to read and write, are not taught to produce or communicate with visuals.

Publication and dissemination. Contemporary scientific communications media are predominantly language-oriented. Printed media are coupled weakly, if at all, to the visual world of space-time. By contrast, half the human neocortex is devoted to processing visual information. In other words, current scientific communication leaves out half — the right half — of the brain. An integral part of our visualization task is to facilitate visual communication from scientist to scientist, engineer to engineer, through visualization-compatible media.

Publication and grants, and therefore tenure, rarely come to researchers whose productivity depends on or produces visualization results. Superiors evaluate scholarly work by counting the number of journal articles published; publications are text, and visual media do not count. Funding itself is based on the careful preparation and evaluation of proposals, which are documents full of words and numbers.

As scientists depend more and more on the electronic network than on the printed page, they will need new technologies to teach, document, and publish their work. Until scientists can build on each other's work, productivity will lag. Publishing (specifically textual materials) has always been a critical part of this building process, and it is one of the primary bottlenecks in CS&E's progress.

Reading and writing were only democratized in the past 100 years. Today, they are the accepted communication tools for scientists and engineers. Table 4 shows that, in time, visualization will also be democratized and embraced by researchers.

Electronic media, such as videotapes, optical disks, and floppy disks, are now necessary for the publication and dissemination of mathematical models, processing algorithms, computer programs, experimental data, and scientific simulations. The reviewer and the reader need to test models, evaluate algorithms, and execute programs themselves, interactively, without an author's assistance. Similarly, scientific publication must extend to use visualization-compatible media.

The use of visualization in scientific computing — in academia, government research laboratories, and industry — will help guarantee

- US preeminence in science and technology,
- a well-educated pool of scientists and engineers with the quality and breadth of experience required to meet the changing needs of science and society, and
- American industries that can successfully compete in the international economic arena.

The information age has yet to deal with information transfer. Visualization technologies can help lead the way to better global understanding and communication. □

References

1. B.H. McCormick, T.A. DeFanti, and M.D. Brown, eds., "Visualization in Scientific Computing," *Computer Graphics*, Vol. 21, No. 6, Nov. 1987.

2. T.A. DeFanti and M.D. Brown, "Scientific Animation Workstations: Creating an Environment for Remote Research, Education, and Communication," *Academic Computing*, Feb. 1989, pp. 10-12, 55-57.

3. T.A. DeFanti and M.D. Brown, "Scientific Animation Workstations," *SuperComputing*, Fall 1988, pp. 10-13.

4. T.A. DeFanti and M.D. Brown, "Insight Through Images," *Unix Review*, Mar. 1989, pp. 42-50.

5. R.A. Drebin, L. Carpenter, and P. Hanrahan, "Volume Rendering," *Computer Graphics* (SIGGraph Conf. Proc.), Vol. 22, No. 4, Aug. 1988, pp. 65-74.

6. C. Upson and M. Keeler, "VBuffer: Visible Volume Rendering," *Computer Graphics* (SIGGraph Conf. Proc.), Vol. 22, No. 4, Aug. 1988, pp. 59-64.

7. G. Campbell et al., "Two-Bit/Pixel Full Color Encoding," *Computer Graphics* (SIGGraph Conf. Proc.), Vol. 20, No. 4, Aug. 1986, pp. 215-223.

Thomas A. DeFanti is a professor of electrical engineering and computer science and codirector of the Electronic Visualization Laboratory at the University of Illinois at Chicago, and an adjunct professor at the National Center for Supercomputing Applications. He is an international lecturer and author in the computer graphics field, and he has had many interactive computer graphics installations in museums and conferences worldwide. He also serves on the editorial boards of several publications, is past chair of ACM SIGGraph, is editor-in-chief of the *SIGGraph Video Review*, was recipient of the 1988 ACM Outstanding Contribution Award, and was recently appointed to the Illinois Governor's Science Advisory Committee.

DeFanti was vice chair of the National Science Foundation-sponsored Panel on Graphics, Image Processing, and Workstations and coeditor of its 1987 report *Visualization in Scientific Computing*. He received his BA in mathematics from Queen's College in New York in 1969 and his PhD in computer science from Ohio State University in 1973.

Readers may contact DeFanti and Brown at the Electronic Visualization Laboratory, Department of Electrical Engineering and Computer Science, University of Illinois at Chicago, PO Box 4348, Chicago, IL 60680.

Maxine D. Brown is associate director of the Electronic Visualization Laboratory at the University of Illinois at Chicago, where she is responsible for funding, documentation, and promotion of its research activities. Brown is an international lecturer and author in the computer graphics field, and she has consulted for a number of companies in the computer graphics industry in the area of professional and technical communications.

Brown is past secretary and vice chair for operations of ACM SIGGraph, is a member of the executive committee of the IEEE Technical Committee on Computer Graphics, and was coeditor of *Visualization in Scientific Computing*. She received her BA in mathematics from Temple University in 1972 and her MSE in computer science from the University of Pennsylvania in 1976. She is a member of the IEEE Computer Society.

Bruce H. McCormick is professor of computer science and director of the Visualization Laboratory at Texas A&M University. McCormick is an international lecturer and author in the areas of computer vision, computer architecture, and scientific visualization. He has held positions at the University of Illinois at Urbana-Champaign, where he was principal investigator of the Illiac III image-processing computer, and at the University of Illinois at Chicago.

McCormick has served on committees of the National Institutes of Health, participated in the 1981 Fifth-Generation Computer Conference in Japan, and was chair of the NSF-sponsored Panel on Graphics, Image Processing, and Workstations and coeditor of *Visualization in Scientific Computing*. He received his BS in physics from the Massachusetts Institute of Technology in 1950 and his PhD in physics from Harvard University in 1955. He is a member of the IEEE.

Acknowledgments

We acknowledge the following individuals for their insightful contributions and comments during the preparation of the NSF report *Visualization in Scientific Computing*: David Arnett, Gordon Bell, James F. Blinn, Frederick P. Brooks Jr., Mel Ciment, James Clark, John Connolly, Jerome Cox, Martin Fischler, Donald Greenberg, Andrew Hanson, Albert Harvey, Laurin Herr, David Hoffman, Robert Langridge, Thomas Lasinski, Carl Machover, Partrick Mantey, Mike McGrath, Nicholas Negroponte, Arthur J. Olson, Jeff Posdamer, Azriel Rosenfeld, David Salzman, Larry Smarr, Alvy Ray Smith, Barna Szabo, Richard Weinberg, Turner Whitted, Karl-Heinz Winkler.

Moving?

PLEASE NOTIFY
US 4 WEEKS
IN ADVANCE

Name (Please Print)

New Address

City State/Country Zip

MAIL TO:
IEEE Service Center
445 Hoes Lane
Piscataway, NJ 08854

ATTACH
LABEL
HERE

- This notice of address change will apply to all IEEE publications to which you subscribe.
- List new address above.
- If you have a question about your subscription, place label here and clip this form to your letter.

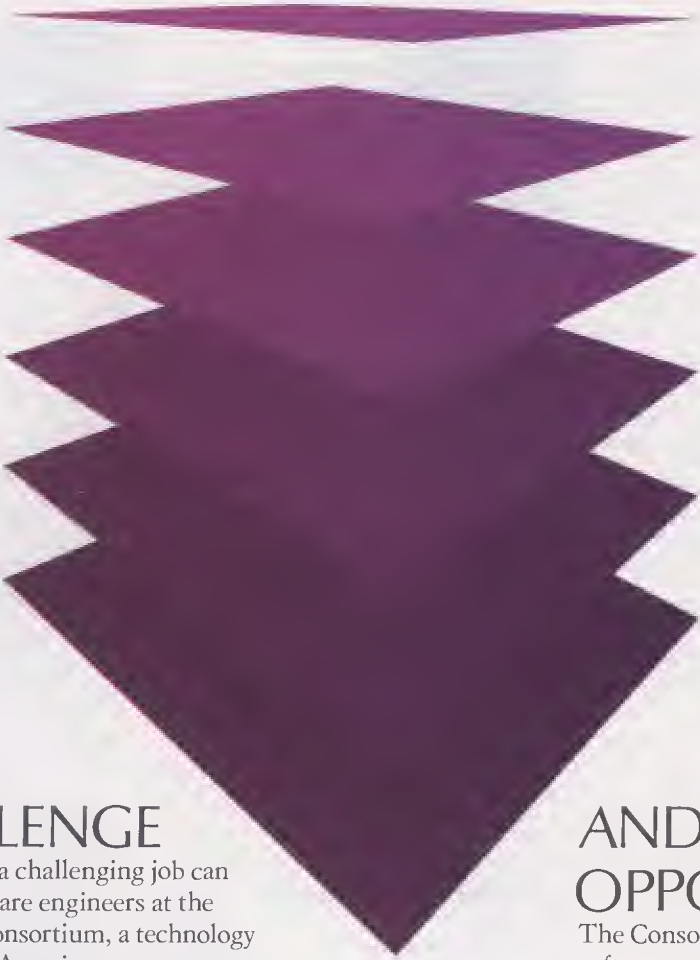

THE CHALLENGE

Professionals know that a challenging job can be its own reward. Software engineers at the Software Productivity Consortium, a technology partnership of 14 major American aerospace companies, are doing things that others in the profession are only talking about. How so? Consortium strategists are building a software development method which incorporates prototyping and reuse and addresses the entire software life cycle. Consortium tool builders are creating products which help implement the method, providing integrated CASE products for work in every development and maintenance phase. They are exploring ways to bridge the gap between system and software engineering, applying a variety of technologies to make model based design and reusable software a reality. Consortium engineers are addressing the technical and management issues posed by programming-in-the-large. They are extending the frontiers of database management technology. They are enriching the concepts of network computing. And they are leading the software industry movement towards tool integration and licensing standards.

Is there a challenge at the Consortium for you?

SOFTWARE PRODUCTIVITY
CONSORTIUM

AND THE OPPORTUNITY

The Consortium is currently seeking software professionals with expertise in one or more of the following areas:

- Database Technology
- Software Reuse
- Software Methodology
- Ada Prototyping, Design, Development
- CASE Tool Development
- Domain Analysis
- Software Prototyping
- Systems Programming
- Computer Systems Administration
- IV&T – Quality Engineering
- Technical Marketing
- User Interface Systems
- Dynamics Assessment
- Technical Education & Training
- System Engineering
- Product Management & Support

SOFTWARE
PRODUCTIVITY
CONSORTIUM ™

SPC Building, Dept. IE89
2214 Rock Hill Road
Herndon, Virginia 22070
Attn: Employment Manager
(703) 742-8877

Representation and Display of Vector Field Topology in Fluid Flow Data Sets

James Helman and Lambertus Hesselink
Stanford University

A representation of the global topology based on an analysis of critical points and their connections is much more readily manipulated and visualized than the original data set.

Visualization of scientific data sets plays an important role in understanding complex phenomena. In fluid mechanics experiments, extensive use of flow visualization has revealed flow structures such as the large-scale eddies in turbulent shear flows discovered by Brown and Roshko[1] in the early 1970s. Their work has had a pronounced influence on fluid mechanics research, because they showed that flow modeling studies could incorporate stable structures in what otherwise appears to be a chaotic flow.

Visualization of physical processes. Traditionally, visualization of complex physical phenomena has been accomplished experimentally as part of the measurement process. For example, extensive studies of fluid flows have relied on seeding the flow with smoke or photographing patterns of oil streaks on the surface of a body in the flow. The resulting data sets, although of very high resolution, are essentially images, well suited to two-dimensional image processing and analysis.

More recently, however, both experiments and numerical simulations have begun to yield high-resolution multivariate data. These incredibly rich and complex data sets consist of several quantities, defined at hundreds of thousands, sometimes millions, of points on a two- or three-dimensional grid. The size of these data sets increases rapidly as the dimensionality grows when the data are dependent on time or on one or more parameters.

For example, a data set with three velocity components, pressure, and density defined on a $100 \times 100 \times 100$ grid defines five quantities at 10^6 points. If sufficient computational power and storage were available, one might want to consider this for 100 time steps and 100 different values of a parameter; the size would jump to 10^{10} points. Even for the more modestly sized two-dimensional data sets we have studied, such as a 60×121 grid at 180 time steps, manipulation and display of the data is computationally very expensive unless some analysis and feature extraction precede the visualization process.

Extensive interaction and feedback between experiment and numerical analysis is necessary both to verify the theoretical models used in the simulations and to tune parameters before embarking on expensive experimental verification of a design. For example, we may wish to improve the efficiency of a physical process, such as combustion. The full characterization of the combustion process requires the measurement of multidimensional data, including density, velocity, and pressure, in a two- or three-dimensional domain. Visualization and interpretation of these data now take on a new aspect, because we want to extract qualitative information and we

0018-9162/89/0800-0027$01.00 ©1989 IEEE

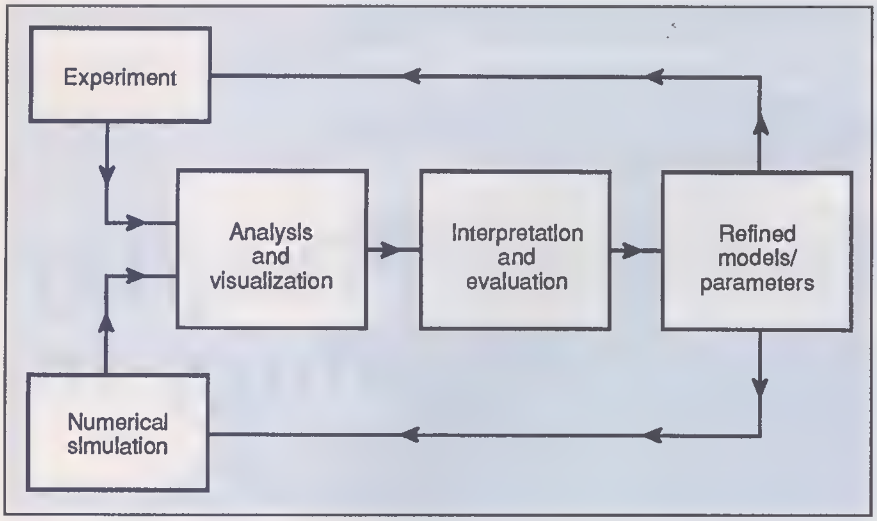
Experiment
Analysis and visualization
Interpretation and evaluation
Refined models/ parameters
Numerical simulation

Figure 1. The role of visualization in simulations and experiments.

need to make quantitative comparisons to iteratively optimize the system.

Ideally, an automated process would use this information to vary input parameters to the experiment or simulation to optimize it for the desired behavior. In most cases of practical significance, however, the processes are insufficiently understood or too complex for such automated optimization. Human reasoning is required, but the data necessary to fully characterize such a complicated physical process are too numerous for human understanding. Therefore, the data must be converted into a form in which humans can comprehend the important aspects. Figure 1 shows this process schematically. Because of the large bandwidth of the human visual system and the speed of computers, the visualization of computer-extracted features provides the most efficient way of effecting this interaction.[2]

Our research effort has concentrated on automated analysis of vector field topology in fluid flow data sets, because topology is central to understanding fluid flows.[3,4] Also, the geometric nature of topology makes it possible to visually communicate this rich and fascinating aspect of fluid dynamics.

The visualization process. Visualization of scientific data includes three processes: analysis, display, and interaction. Most large multivariate data sets are not suited to direct display (imagine 800,000 little color-coded arrows in space!), so analysis is almost always required. Unless the display is founded on an analysis or interpretation relevant to the problem, you end up with a pretty picture that has little scientific value.

Analysis. The analyses we can apply to multivariate data range from general purpose techniques, such as contour surface extraction, to highly specific ones, such as shock surface or vortex core extraction in the study of fluid flow data sets. The method described in this article lies between these extremes. One can apply the analysis of critical points and topology to any vector field, but we have adapted the analysis somewhat to situations occurring in fluid flows by including in the analysis

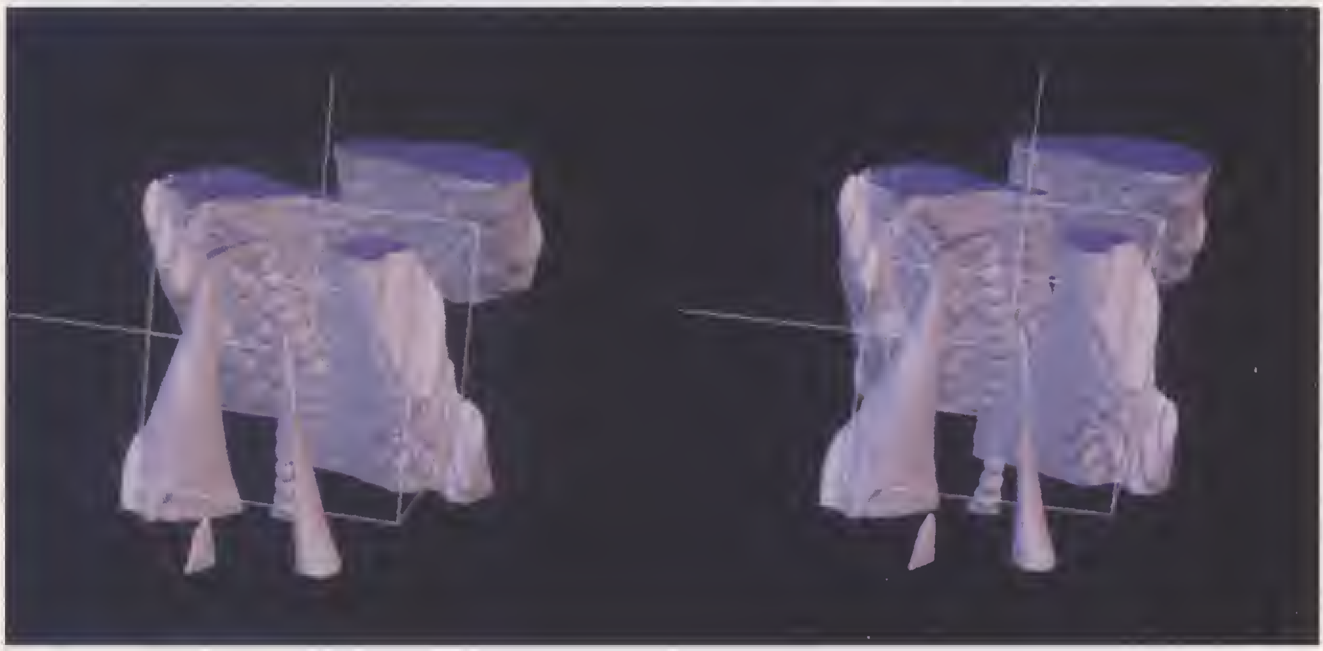

Figure 2. Isosurface of the velocity magnitude in the experimentally measured flow around a circular cylinder (shown as a cross-fusable stereo pair). See the sidebar for viewing instructions.

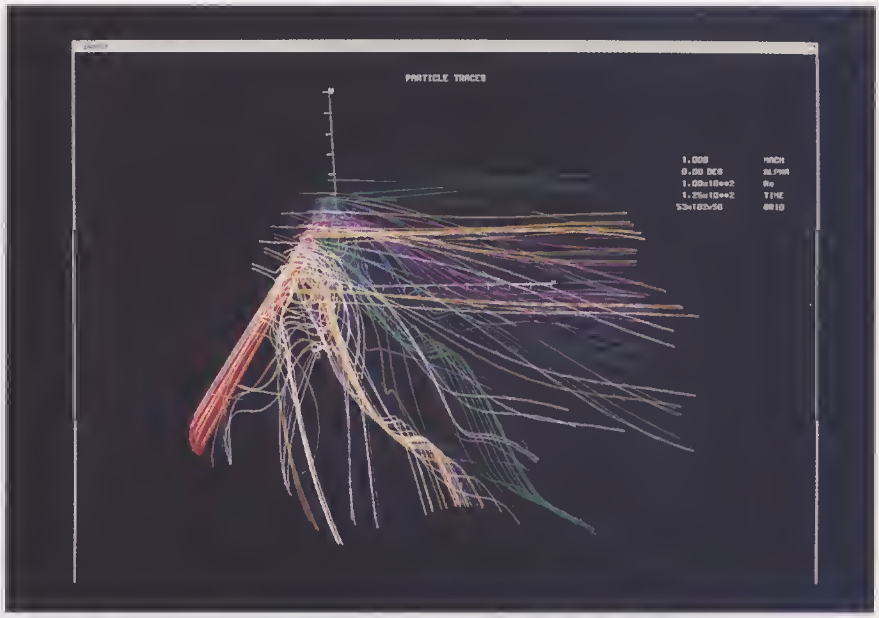

Figure 3. Selected velocity field tangent curves in the computed flow around an airfoil at 90 degrees.

walls on which the magnitude of the vector vanishes.

Display. The more complex the data set, the more important it is to maximize the amount of information understandable from the display. This requires tailoring the display to match the capabilities of the human visual system. This tailoring includes the use of colors, surfaces instead of isolated curves, orientation cues such as the lighting and shading of surfaces, depth cues such as stereopsis, and animation.

Interaction. Even with the best display techniques, it is impossible to simultaneously display all of the quantitative information in a large data set in an understandable form. So, when we require more precise information on some element or aspect of the data, we must specify it through user interaction. This interaction can include simple refinement of the color coding of contour levels or complex manipulation of virtual probes that simulate physical behavior or measurements.

Vector field visualization. In the past, several graphics display methods were developed for scalar data sets, but these methods often cannot help us display and understand vector data sets. The difficulty arises from the limitations of our vision and display devices. We cannot directly display vectorial data on a two-dimensional screen, for example as a set of little arrows, and still interpret the result with the same ease as we would a scalar image. Our visual systems simply are not well adapted to interpret large volumes of vectors in this way, whereas we have superb abilities for understanding and interpreting images or depth-cued surface displays.

One method of generating surfaces is to apply a scalar display technique such as isovalue surface extraction to a scalar quantity derived from the vector field. For example, Figure 2 shows an isovalue surface of the velocity magnitude for the experimentally measured two-dimensional flow around a circular cylinder[5] as a function of time. The surface was generated by StanSurfs,[6] which implements a spline interpolation and subdivision technique for extracting contour surfaces.[7] The figure does not show the cylinder, which is located in the foreground. Time increases from top to bottom, and the flow within each time slice is directed into the page. At each instant, each contour line surrounds a vortex, a saddle, or a vortex-saddle pair. You can see the time development of the flow downstream of the cylinder in the separation and merging of the contour surfaces from top to bottom.

Unfortunately, we lose a great deal of information when we reduce a vector quantity to a scalar for visualization. Experimental visualization techniques, such as smoke seeding, often yield information about particle paths or instantaneous streamlines. Thus, for the evaluation of simulations and their verification against experiment, flow visualization packages, such as Plot3D,[8] usually include the ability to visualize tangent curves of velocity and vorticity.

These curves prove extremely valuable in understanding fluid flows, but when we display more than a few well-grouped curves, the display often becomes difficult to understand because the isolated curves are poorly suited to depth cuing and visual interpretation, as shown in Figure 3. Ani-

Stereo viewing instructions

The stereo pairs in this article are presented for cross-fused viewing with the left eye's view on the right and the right eye's view on the left. The pairs are most easily viewed using special goggles with mirrors, but most people can view them directly with a little patience and practice. The pairs are best viewed at a distance between 16 and 32 inches. Closer viewing is possible, but more stressful because it requires you to cross your eyes more.

Cross your eyes until you see four objects, then slowly reduce the amount of crossing until the middle two images fuse into one. The merged image is usually blurred at first, but you can bring it into focus by slowly relaxing your eyes' focus to infinity while keeping the two middle images merged.

Some people find that holding an index finger down along the bridge of the nose or moving it between the eyes and the page helps achieve the proper merging and focusing. It also helps if you frame the images with black paper. If one image appears higher on the page than the other, rotate the page until they merge. It takes some practice, but once the images are merged and in focus, you can easily maintain the viewing.

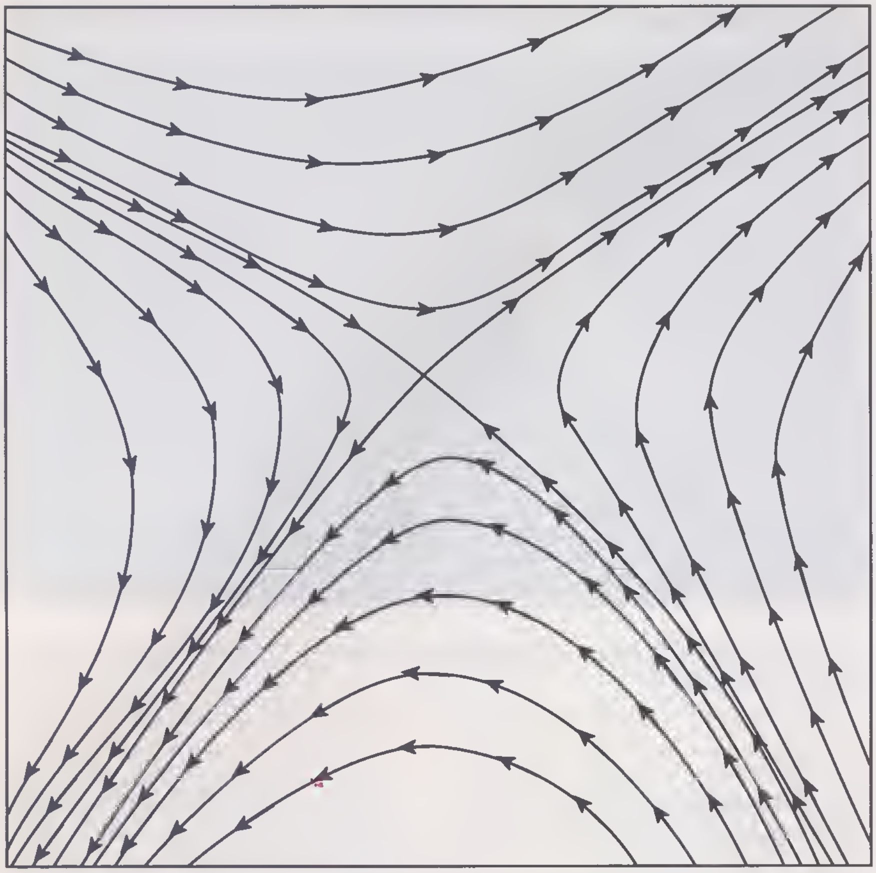

Figure 4. Vector field and tangent curves near a saddle point.

mation reduces the problem by displaying only a small section of the tangent curves at any one time,[9] but there remains the problem of selecting starting points for the curves.

The usual method for determining and visualizing flow topology in a complex flow data set is to chose points at which to start the integration of tangent curves and to interactively refine the starting points until the desired structures are located. The researcher then hand draws surfaces to represent the topology.

We are investigating an alternative approach based on the physics and mathematics underlying the physical phenomenon. In fluid flows, we can determine (assuming a particular frame of reference) critical points in the flow. These are points where the velocity vector vanishes. The critical points, connected by principal lines or planes, determine the topology of the flow.

The decomposition of the flow field into elementary structures provides an effective simplification. Not only have we reduced the complexity of the data set, but we have achieved this without sacrificing the quantitative nature of the data set. The positions of the structures and surfaces are accurately contained in the representation. From this representation, we may directly visualize the topology as surfaces or use the internal graph representation to compare topologies and locate topological transitions.

Analysis

We can think of flow topology in terms of surfaces in three-dimensional flows or curves in two-dimensional flows that divide the flow into separate regions. Two sets of surfaces or curves are of particular interest.[10] The first set consists of those along which the flow close to the wall of a body attaches to or separates from that wall, that is, those tangent curves that actually end on the wall rather than moving along the surface. The second set consists of those surfaces where tangent curves that start arbitrarily close to each other can end up in substantially different regions. This second group of curves relates to critical points. For example, tangent curves starting on either side of a curve that goes directly into a saddle point are diverted by the saddle point to very different regions.

We have developed algorithms for generating representations of the topology for two-dimensional velocity fields. When the two-dimensional field is available for a discretely sampled range of times or parameter values, we link these two-dimensional representations together to form a three-dimensional representation of the flow with time or the parameter as the third dimension.

Tangent curves. A tangent curve of a vector field is a curve for which the tangent vector at any point along the curve is parallel to the vector field at that point. The tangent vector at each point along the curve is the derivative of the position vector along the curve with respect to arc length.

Figure 4 shows several tangent curves in the neighborhood of a saddle point and some of the vectors that generated them. We derive tangent curves by integrating the vector field. Both the magnitude and the direction of the vector field are determined by the tangent curves when the curves are parameterized by the integration parameter instead of arc length.

For a velocity field, the curve parameter is the travel time along the trajectory. If enough tangent curves are known, the entire vector field can be reconstructed from them.

The power of topological visualization is that, given the critical points and their principal tangent curves, an observer can visually infer the shape of other tangent curves and hence the structure of the whole vector field.

Critical points. Critical points are those points at which the magnitude of the vector vanishes. We can characterize these points according to the behavior of nearby tangent curves. If a tangent curve integrates directly into one of these points, the integration must stop, since the tangent is undefined at this point. The tangent curves that end on critical points hold special interest because they define the behavior of the vector field in the neighborhood of the point.

A critical point can be classified to a first-order approximation according to the eigenvalues of the Jacobian matrix of the vector (u,v) with respect to position at the critical point $(x0, y0)$:

$$\left.\frac{\partial(u,v)}{\partial(x,y)}\right|_{x0,y0} = \left.\begin{bmatrix} \frac{\partial u}{\partial x} & \frac{\partial u}{\partial y} \\ \frac{\partial v}{\partial x} & \frac{\partial v}{\partial y} \end{bmatrix}\right|_{x0,y0}$$

Figure 5 shows how the eigenvalues classify a critical point as an *attracting node*, a *repelling node*, an *attracting focus*, a *repelling focus*, a *center*, or a *saddle*. We can understand this by observing that a positive or negative real part of an eigenvalue indicates an attracting or repelling nature, respectively. The imaginary part denotes circulation about the point. Among these points, the saddle points are distinct in that only four tangent curves actually end at the point itself. At the saddle point, these curves are tangent to the two eigenvectors of the Jacobian matrix, which are the separatrices of the saddle point. The outgoing and incoming separatrices are parallel to the eigenvectors with positive and negative eigenvalues, respectively.

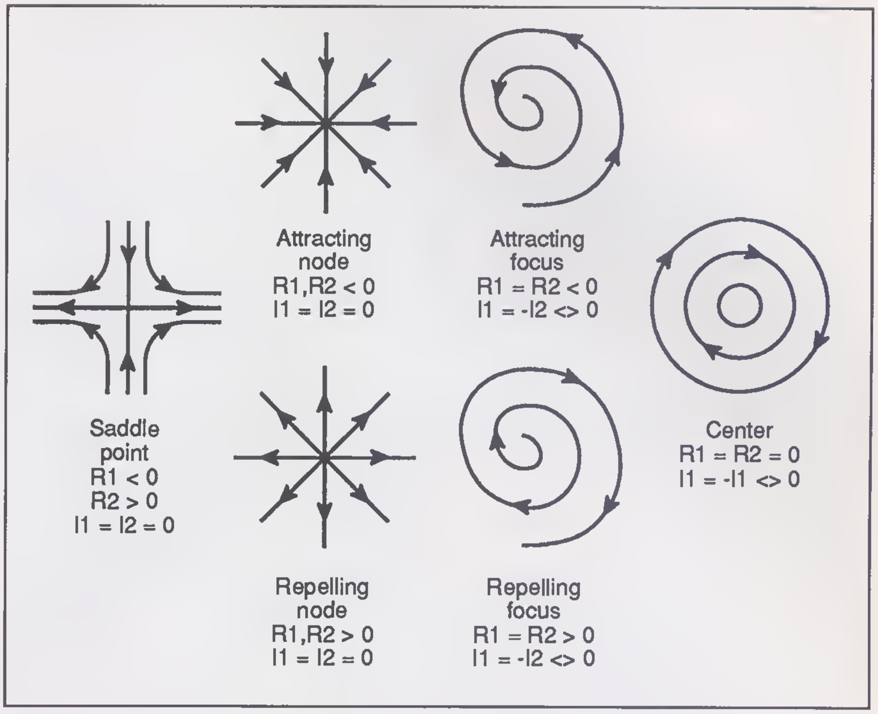

Figure 5. Classification criteria for critical points. R1 and R2 denote the real parts of the eigenvalues of the Jacobian; I1 and I2 denote the imaginary parts.

On walls where the velocity is constrained to be zero, the normal component of the velocity usually vanishes as the square of distance from the wall, and the tangential component vanishes linearly. This causes tangent curves very close to the wall surface to remain close to it. However, there can be singular points at which a tangent curve terminates on the body rather than moving along the surface of it. These points are *attachment nodes* and *detachment nodes*.

The saddle points and attachment and detachment nodes are the only points where unique tangent curves end on the point itself. Because these points serve as the origin of the tangent curves that characterize the topology, we call them *originators*. All other points are called *terminators*.

Some of the tangent curves end at other critical or wall points. Others may leave the domain of available data. The points at which these tangent curves leave the domain we call *boundary incoming* and *boundary outgoing* points, depending on whether the vector was parallel or antiparallel to the direction of integration when it left the domain.

Based on this, we classify each point as a saddle point, attracting node/focus, repelling node/focus, center, attachment node, detachment node, boundary incoming, or boundary outgoing. These classifications imply membership in other classes as shown in Figure 6. For simplicity, we excluded centers from the diagram and the subsequent discussion because no tangent curves enter or leave these points. The outer closed tangent curve surrounding a center is represented by the tangent curves originating from the saddle, which by virtue of continuity and boundary conditions must accompany it. Therefore, this entails no loss of global structural information.

Algorithm for two-dimensional flows. Starting from a two-dimensional vector field defined on a connected grid of points, which is specified by a separate array, we locate and characterize the critical and wall points. Each of these points is represented by a data object that has slots containing intrinsic information about the point, including position, the Jacobian matrix, eigenvalues, and eigenvectors. For each point, these quantities are computed and stored both relative to the grid and relative to physical space. Slots are also provided, as appropriate to the class of the object, to link incoming and outgoing tangent curves.

For those points having principal directions associated with them, namely the *originators*, tangent curves are integrated forward along the outgoing directions and backward along the incoming directions. These tangent curves are used to link the points into a partially connected graph.

Given a two-dimensional vector field, the algorithm is as follows:

(1) Locate, characterize, and classify all critical and wall points.

(2) Integrate tangent curves out of the originators. This generates one curve for each attachment or detachment point and four curves for each saddle point. The beginning of each curve is linked to its originator.

(3) Identify the endpoint of each tangent curve and link the curve to it.

(a) If the curve ends on the boundary, a new boundary point is created to represent it, and the curve is linked to it.

(b) If the curve ends at an existing terminator, then the end of the curve is linked to it.

(c) If the endpoint is an originator, then the corresponding tangent curve for that originator is replaced by the new curve. This situation occurs because of integration errors. This algorithm gives preference to the topologically unstable connections between originators.

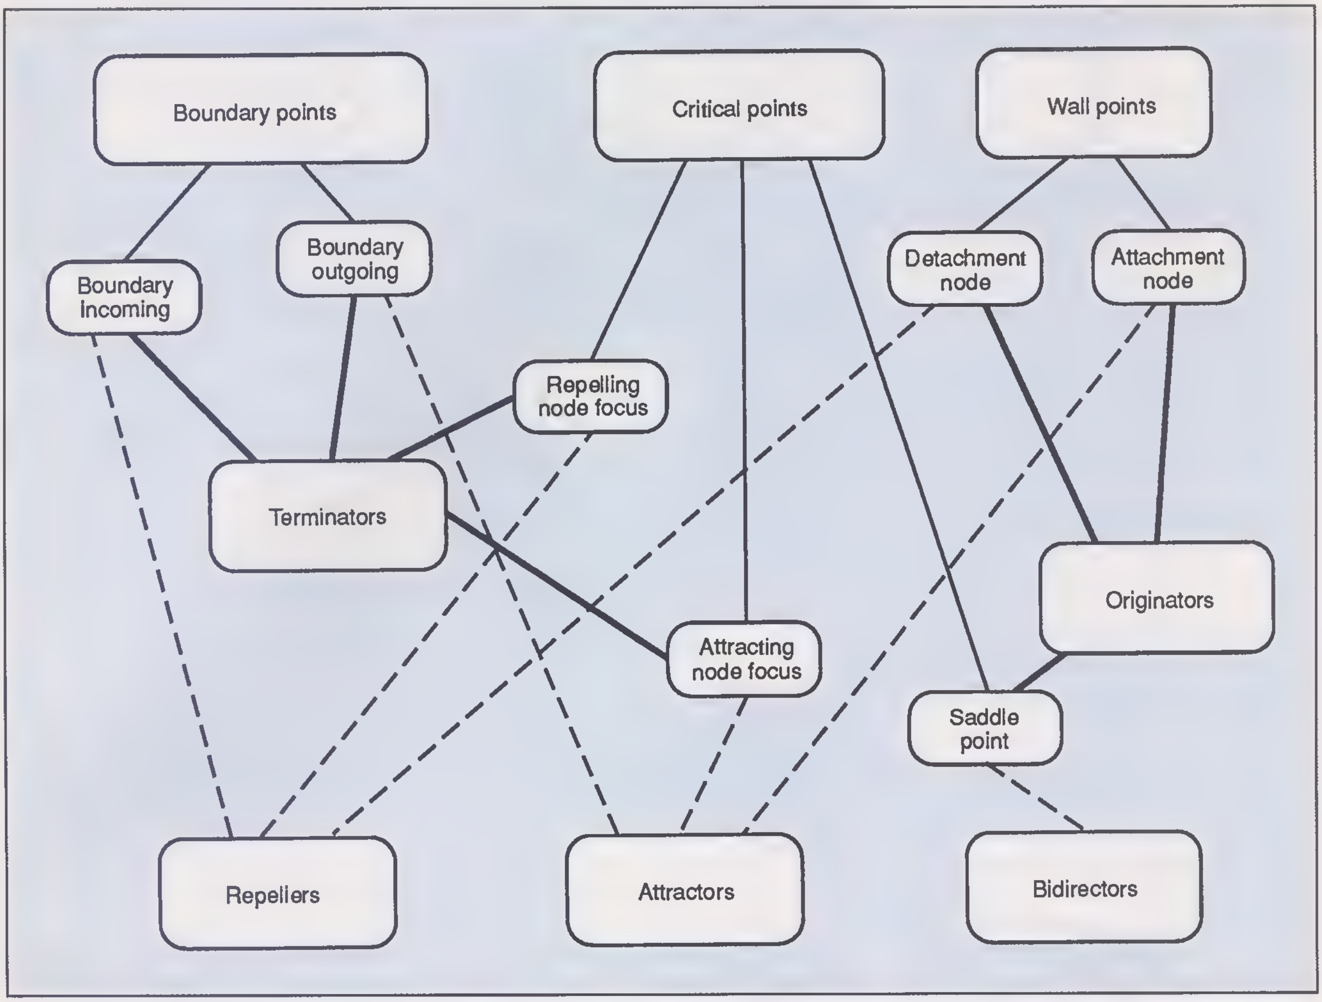

Figure 6. Point classifications. Each point identified by the analysis belongs to one of the classes denoted by the small boxes. For each of these point classes, this implies membership in three superclasses as shown by the connecting lines.

Figure 7a shows the results of the application of this algorithm to the computed two-dimensional flow around an airfoil at 90 degrees.[11] A saddle point, labeled *sp*, and two foci, labeled *af* and *rf*, behind the airfoil indicate the locations of two vortex cores. These points, as well as the boundary incoming (*bi*) and outgoing (*bo*) points appear in a schematic graph representation in Figure 7b. This format helps us compare topologies because it "unwinds" the curves and unambiguously portrays the topology.

This particular topology is unstable because two originators, a saddle point and a detachment node, labeled *de*, are joined by a tangent curve. Normally, this configuration only exists for one instant in time, but the instant can be spread out slightly by numerical errors. At a slightly later time, the topology has changed to that shown in Figure 8.

Algorithm for two-dimensional parameter-dependent flows. When the two-dimensional vector field depends on some parameter such as time, we can link the representations of instantaneous topologies to form a three-dimensional representation. The order in which we consider points and curves in different slices for linking is important because the distance in the plane between two points of the same class is not sufficient to indicate that they have the same place in the flow topology. This is particularly true of the boundary incoming and boundary outgoing points, which can be very close together both with respect to the grid and in physical space. For this reason, we use the originators as a basis for the linkage. Points connected to corresponding curves from an originator are linked recursively to ensure consistency.

We link slices by applying the following method to the representation for each parameter value:

For each subclass of originators and then for each subclass of terminators, do the following:

(1) Construct the set of all parts consisting of points of that subclass with the first element from the current slice and the second element from the next slice.

(2) Keep only those pairs in which neither of the points is already linked to a point in the other slice and for which the two-dimensional distance separating them is less than a specified tolerance.

(3) Sort these pairs in ascending order according to the separation distance between the points in the pair.

(4) For each pair of points, starting with the ones with the smallest separation:

(a) Link the points together.

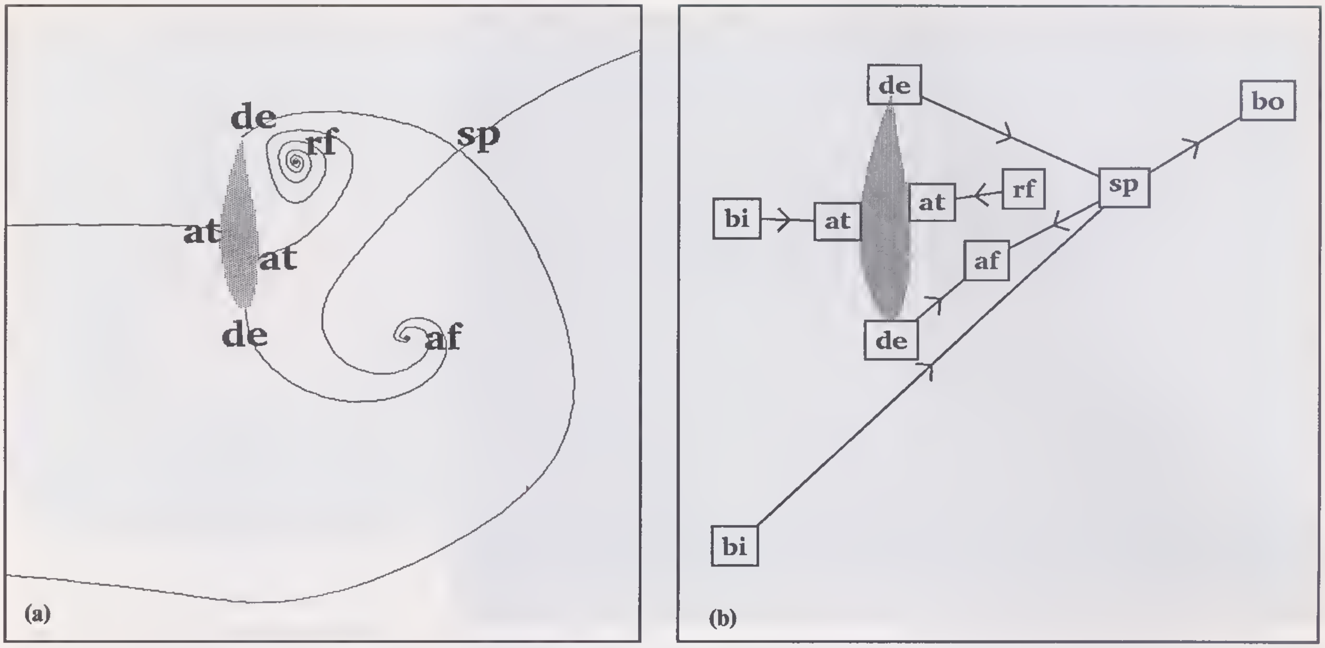

Figure 7. Tangent curves (instantaneous streamlines) and points in the computed flow around an airfoil at t=2 (a). Corresponding graph of flow topology (b).

(b) If there are multiple tangent curves from the point, align them according to direction.

(c) For each pair of tangent curves, if the points at the other end of the tangent curves satisfy the criteria in (2), link the points together by recursing to (a).

Display

For two-dimensional parameter-dependent vector fields, the representation of the topology can be displayed as a set of surfaces with the third dimension corresponding to the parameter value. The surfaces are created by filling in strips between corresponding tangent curves in each slice of the representation. These strips are drawn only when the starting points and endpoints of the first curves are linked to the starting points and endpoints of the second. When a topological transition occurs, one of the endpoints changes, and no strip is drawn. Figure 9 shows the strips generated by the t=3 and t=4 slices. This corresponds to the topology shown in Figure 7.

In the display of these surfaces, we use several cues to aid visualization. The surfaces are lighted and shaded, and they are colored according to their type. Surfaces

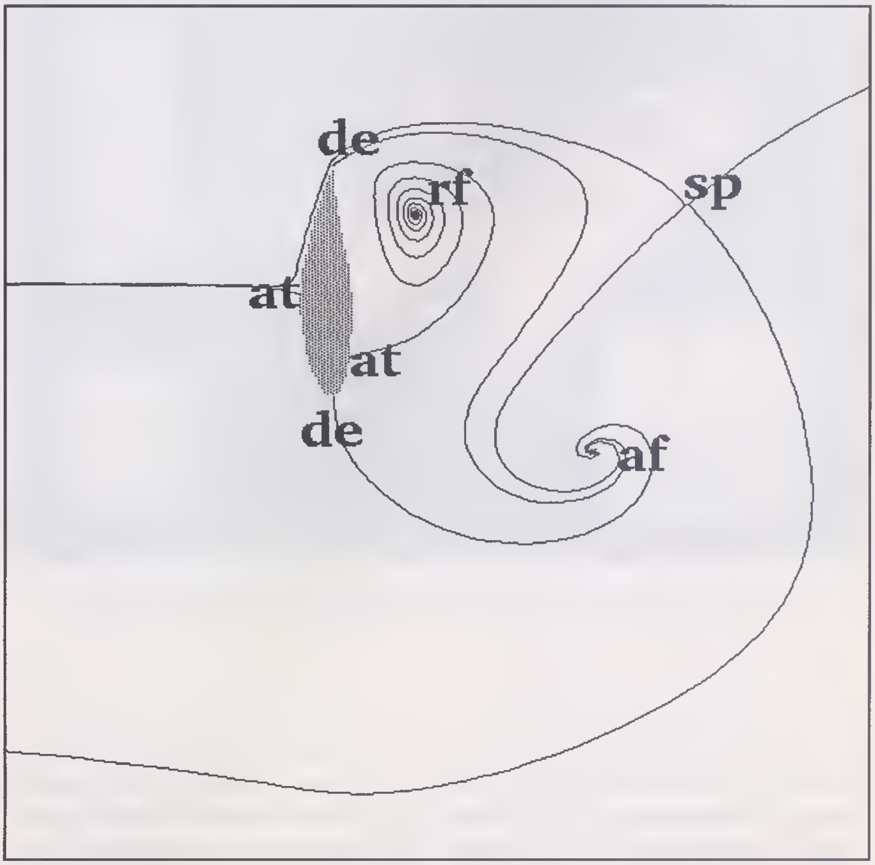

Figure 8. Tangent curves (instantaneous streamlines) and points in the computed flow around an airfoil at t=3.

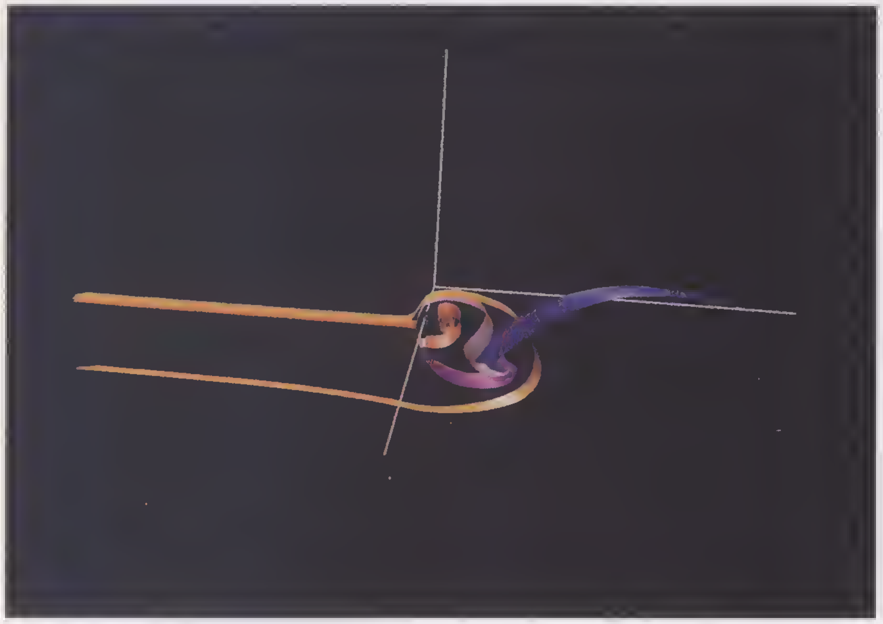

Figure 9. Time slice of the topological surfaces in the computed flow around an airfoil at 90 degrees.

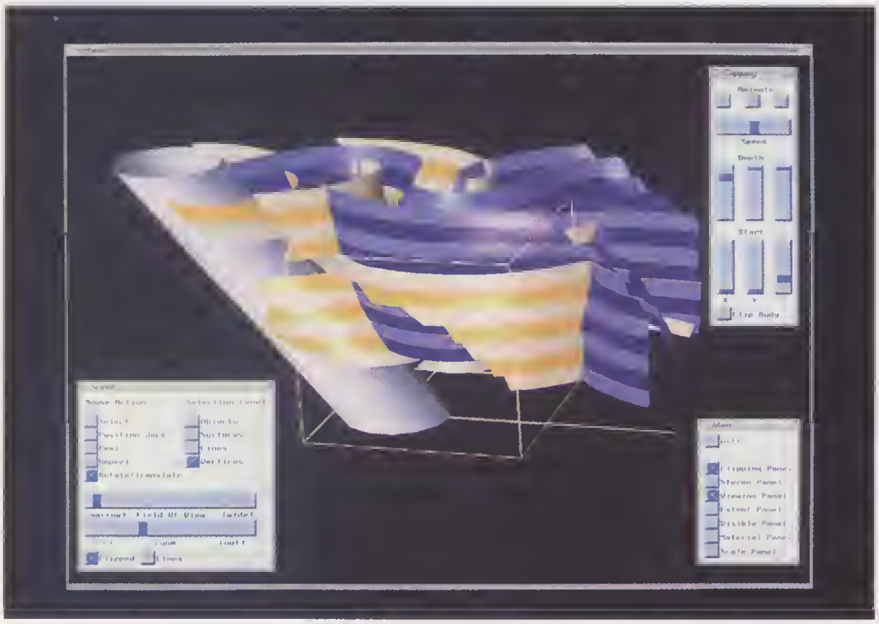

Figure 10. Topological surfaces in the experimentally measured flow around a circular cylinder.

corresponding to the attracting direction of a saddle point are colored yellow. Those surfaces corresponding to the repelling direction are colored blue. Surfaces of attachment are colored orange. Surfaces of detachment are colored purple.

To aid in distinguishing the edges of surfaces and to provide more shape cues, we candy-stripe the surfaces parallel to the slices. The superimposition of a wireframe cube and stereo viewing provide additional depth cues.

Figure 10 shows the surfaces in the experimentally measured cylinder flow shown in Figure 2. Time increases from bottom to top. Parts of four saddle points are shown: two complete ones in the middle of the flow, one forming at the top of the cylinder, and one in the foreground on the far right. Near the top of the latter, a change in topology is visible as an outgoing (blue) surface from one of the upstream saddles moves out from behind the incoming (yellow) surface of the far right saddle.

Figure 11 shows the surfaces in the airfoil flow for a few cycles of this periodic flow. Time increases from back to front. The orange surface coming in from the left terminates on the body of the airfoil and separates those instantaneous streamlines that go over the top edge of the airfoil from those that pass beneath it. A repelling focus behind the airfoil is wrapping a surface of detachment with the outgoing surface from the saddle point in the lower foreground.

Interaction

When a complex three-dimensional scene is projected on a two-dimensional computer screen, we cannot perceive the complete structure of the scene without some interaction. The manipulations common to most interactive display systems include interactive zoom, rotation, and translation, but these are insufficient for the examination of structures occluded by surfaces on all sides. Some of these hidden structures can be made visible by adding x, y, and z clipping planes in the object frame.

However, these generic methods still don't allow sufficient control of the view. A better technique would take advantage of the natural hierarchy of objects in the display: some objects consist of the set of surfaces attached to a particular stack of originator points; each surface consists of stacks of strips; and strips are made up of a set of quadrilaterals. Given this hierarchy, we have included the capacity to peel structures of any level: objects, surfaces, strips, and quadrilaterals. Peeled structures are made invisible and pushed onto a stack so they can be unpeeled on request.

One problem with graphics displays is the frequent difficulty in getting quantitative information from them. For this reason, we decided to add a visible jack, which can be moved around on the objects and prints out its position on the screen. The jack can be seen near the repelling focus structure in the foreground of Figure 11.

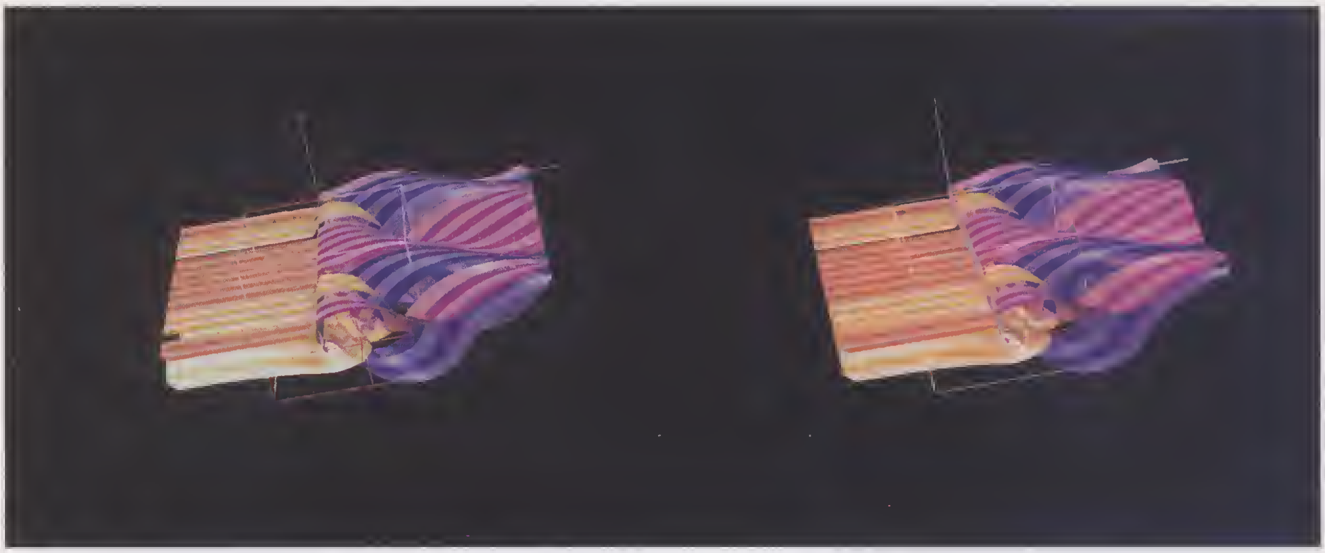

Figure 11. Topological surfaces in the computed flow around an airfoil at 90 degrees (shown as a cross-fusable stereo pair). See the sidebar for viewing instructions.

Lighting and shading parameters are also important to viewing. In particular, specular reflections from surfaces give valuable orientation cues, but depend heavily on the position of the light sources. Hence, the positions and colors of the lights, as well as the colors and characteristics of the materials, are under interactive control.

Implementation

As implemented, the software consists of two parts, the analysis package and the display package.

Analysis workbench. The analysis package is written in Common Lisp using the Portable Common Loops (PCL) implementation of the Common Lisp Object System. This combination has allowed us to develop a workbench on which we can load in various data sets, analyze parts of them, and display the results. Since methods can be interactively applied to the data objects using the interpreter, both the portion of the data under analysis and the algorithms to be applied can be quickly modified and the data reanalyzed without having to reload and restart the analysis from the beginning. For these reasons, Lisp's ability to mix interpreted and incrementally compiled code, as well as its garbage collection of discarded objects, greatly facilitated the development of the algorithms.

We accepted some penalty in floating-point performance for using Lisp rather than a compiled and linked language such as C. In spite of the large number of computations required by the analysis, we have found the increased ease of development more than compensates for the somewhat lower performance. For the airfoil flow with a 53 × 102 grid at 50 time steps, the analysis of this data set takes less than 45 minutes on a Sun-3/260 workstation with 24 megabytes of RAM and a floating-point accelerator board.

One concession to floating-point performance was the use of the tangent curve integration routines from Plot3D. These Fortran routines, called from Lisp via a foreign function interface, allow us to use the floating-point accelerator, which is not directly supported by the Lisp we used.

The analysis package includes a small, object-oriented, two-dimensional graphics system written to display the intermediate two-dimensional results (see Figures 4 and 5). The graphics are created using CLX, the Common Lisp interface to X Windows.

The data sets are read in using the single-grid, three-dimensional data format used by Plot3D. This format defines separate field and grid arrays. Hence, for each grid point five quantities are defined: the x, y, and t coordinates of the grid point and the x and y velocity components. The single-precision data set (53 × 102 × 50) occupies about 5.5 megabytes of storage.

Display program. The display program is written in C using Silicon Graphics' Graphics Library. Rotation of the objects is accomplished with the mouse by clicking and dragging on an edge of a wireframe cube, which causes the object to rotate around the coordinate axis parallel to the edge so that the edge follows the cursor.[12] Translation is specified by clicking and dragging outside of the cube.

Almost every other aspect of viewing can be controlled by the mouse using buttons and sliders created using the Panel Library.[13] Sliders interactively control zoom, field of view, light positions and colors, material colors and characteristics, stereo spatial and angular separation, and the x, y, and z clipping planes. Other capabilities include using the mouse to select, peel, or unpeel at the four levels of the hierarchy: objects, surfaces, strips, or quadrilaterals. The jack mentioned before can be dragged around on surfaces; it causes the corresponding x, y, and z coordinates to be displayed. Some of the control panels can be seen in Figure 4.

The surfaces are drawn from a triangular mesh using Gouraud shading. The surfaces for the 53 × 102 × 50 airfoil data set are defined by approximately 34,000 vertices, most of which are drawn twice since most curves are the boundaries of two strips. This number is small enough to allow for interactive drawing speeds on the workstation used for the display, a Silicon Graphics Iris-4D/80GT.

Interpretation of large, multidimensional vector fields is a difficult task, and presently available techniques are not adequate. The difficulty stems from the inability of the human visual system to assimilate displays containing a large number of vectors or curves. On the other hand, the visual system is superbly adapted to surface displays.

By reducing the original vector field to a set of critical points and their connections, we have arrived at a representation of the topology of a two-dimensional vector field, which is much smaller than the original data set but retains with full precision the information pertinent to the flow topology. This representation can be displayed as a set of points and tangent curves or as a graph, which is especially useful for comparing data sets and detecting topological transitions. When time defines a third dimension, the representation can be readily displayed as surfaces.

We plan to apply the two-dimensional techniques described in this article to the skin-friction vector field near the surface of a body in a three-dimensional flow. This will provide a valuable basis for comparison of computational results with skin-friction visualizations generated experimentally by photographing oil as it streaks along the surface of a body. Starting from a knowledge of the surface topology, we will then be able to analyze and characterize structures in the three-dimensional volume surrounding the body and generate complete representations of the topology of three-dimensional flows. ☐

Acknowledgments

We wish to thank Brian Cantwell of Stanford University and Stuart Rogers of NASA Ames for providing us with interesting data sets. We also wish to thank Malcolm Slaney for providing evidence that Lisp could be used for number crunching.

Numerous free software packages have greatly aided our work. These include Plot3D from Pieter Buning of NASA Ames, the Panel Library from David Tristram of NASA Ames, PCL from Gregor Kiczales of Xerox PARC, Boxview from Kelvin Thompson of the University of Texas, the X11 windowing system from MIT, and gnuemacs from the Free Software Foundation.

This work is supported by NASA under contract NAG-2-489-S1, including support from the NASA Ames Numerical Aerodynamics Simulation Program and the NASA Ames Fluid Dynamics Division.

References

1. G.L. Brown and A. Roshko, "On Density Effect and Large Structure in Turbulent Mixing Layers," *J. Fluid Mechanics*, Vol. 64, pp. 775-816, 1974.
2. L. Hesselink, "Digital Image Processing in Flow Visualization," *Ann. Review of Fluid Mechanics*, Vol. 20, pp. 421-485, 1988.
3. M. Tobak and D. Peake, "Topology of Three-Dimensional Separated Flows," *Ann. Review of Fluid Mechanics*, Vol. 14, pp. 61-85, 1982.
4. U. Dallmann, "Topological Structures of Three-Dimensional Flow Separation," Tech. Rep. DFVLR-IB 221-82 A 07, Deutsche Forschungs und Versuchsanstalt fur Luft und Raumfahrt, Apr. 1983.
5. B. Cantwell and D. Coles, "An Experimental Study of Entrainment and Transport in the Turbulent Near Wake of a Circular Cylinder," *J. Fluid Mechanics*, Vol. 136, pp. 321-374, 1983.
6. J. Helman et al., "StanSurfs: A Contour Surface Reconstruction and Display Package," Fourier Optics and Optical Diagnostics Laboratory, Durand Building, Stanford University, Stanford, Calif. Copies of the document and software are available from L. Hesselink.
7. K. Wu and L. Hesselink, "Computer Display of Reconstructed 3D Scalar Data," *Applied Optics*, Vol. 27, Jan. 1988, pp. 395-404.
8. P. Buning and J. Steger, "Graphics and Flow Visualization in Computational Fluid Dynamics," in *AIAA 7th Computational Fluid Dynamics Conf.: A Collection of Technical Papers*, American Institute of Aeronautics and Astronautics, June 1985, Paper 85-1507-CP.
9. T. Lasinski et al., "Flow Visualization of cfd Using Graphics Workstations," in *AIAA 8th Computational Fluid Dynamics Conf.: A Collection of Technical Papers*, American Institute of Aeronautics and Astronautics, June 1987, Paper 87-1180, pp. 814-824.
10. J. Helman and L. Hesselink, "Automated Analysis of Fluid Flow Topology," in *Three-Dimensional Visualization and Display Technologies, SPIE Proceedings*, Vol. 1,083, Jan. 1989, Paper 1083-23.
11. S. Rogers and D. Kwak, "An Upwind Differencing Scheme for Time-Accurate Incompressible Navier-Stokes Equations," in *AIAA 6th Applied Aerodynamics Conf.: A Collection of Technical Papers*, American Institute of Aeronautics and Astronautics, June 1988, Paper 88-2583.
12. K. Thompson, "Boxview, an Interactive Graphics Program," University of Texas, posted to Unenet newsgroup comp.sys.sgi, July 1988.
13. D.A. Tristram and P.P. Walatka, "Panel Library: Programmer's Manual, Version 9," tech. report, NAS Systems Div., NASA Ames Research Center, Moffett Field, Calif., March 1989.

James L. Helman is a research assistant in the Fourier Optics and Optical Diagnostics Laboratory at Stanford University, where he is pursuing a PhD in applied physics. His dissertation work is on the interpretation and visualization of large data sets. Other research interests include computer graphics, dynamical systems, and scientific computing. Before beginning his thesis project, he worked as a consultant writing system-level software for peripheral image and array processors.

Helman received a BA in physics and math from Washington University, St. Louis, in 1982 and an MS in applied physics from Stanford in 1984. He received an A.H. Compton science fellowship from Washington University and a graduate fellowship from the National Science Foundation. He is a student member of the APS, ACM, and IEEE Computer Society.

Lambertus Hesselink holds a joint appointment as associate professor in the Aeronautics and Astronautics Department and the Electrical Engineering Department at Stanford University. His research interests include nonlinear optics, optical phase conjugation, optical signal processing, optical diagnostics, and 3D digital image processing. Before joining the faculty at Stanford, he was an instructor in applied physics at the California Institute of Technology.

Hesselink received a BS in applied mechanics in 1970 and a BS in applied physics in 1971, both from the Twente Institute of Technology in the Netherlands. He received an MS in 1972 and a PhD in 1977 from the California Institute of Technology.

He was awarded the Stheeman prize in 1970, was elected the AIAA teacher of Northern California in 1984, and was a corecipient of the Itek prize in 1985. Hesselink is listed in numerous *Who's Who* and is a member of the AIAA, APS, OSA, SPIE, and ACM.

Readers may contact the authors at the Department of Aeronautics and Astronautics, Durand Building Room 370, Stanford University, Stanford, CA 94305-4035.

The joy of C-scape

The C-scape™ Interface Management System frees C programmers from the tedium of coding windows, menus, data validation, help, and text editing functions.

Moreover, C-scape is a joy to use. With C-scape's object-oriented design, you'll build more functional, more flexible, more portable, and more unique applications—and you'll have more fun doing it.

The industry standout. Many thousands of programmers have quit home-grown libraries and cumbersome, inflexible products for the pleasure of C-scape. The press agrees: "C-scape is by far the best . . . a joy to use," wrote IEEE *Computer. PC Magazine* chose C-scape to produce its Laboratory Benchmark Series 5.0 software because C-scape offers mouse support. Moreover, C-scape simultaneously combines text and graphics. And because C-scape makes it easy to create your own custom routines, major companies have selected C-scape as a standard for software development.

C-scape is built around an open architecture, so you can use it with data base management or other C libraries. And to port from MS-DOS or OS/2 to UNIX, just recompile.

Trial with a smile. C-scape is not only the most sophisticated, flexible and powerful interface system available, it's also the most friendly—and easiest to use. Try C-scape on a 30-day trial. It comes with a thorough manual, demo disk, sample programs with source code, an optional screen designer and code generator, access to a 24-hour bulletin board, and toll-free support. No royalties, runtime licenses, or runtime modules. After you register, you get complete library source code at no extra cost.

Call 800-233-3733 (617-491-7311 in Mass.) to try C-scape now. After the joy of C-scape, programming wil never be the same.

MS-DOS, OS/2: $299, library only; with Look & Feel, $399. UNIX, XENIX, Apollo, Sun, Stratus, others please call. Mastercard and Visa accepted.

C-scape Features

Graphics. Combine high-resolution color graphics with text or menus.

Object-oriented. Add features and create reusable code modules.

Mouse. Use any standard mouse for fast screen control.

Portability. Write hardware independent code. Supports DOS, OS/2, UNIX, others. Autodetects Hercules, CGA, EGA, VGA.

Text editing. Create a full-featured text editor or pop-up note pad.

Field flexibility. Create masked, protected and marked fields with complete data validation. Use time, date, money, pop-up list, and many more functions, or create your own.

Windows. Choose from pop-up, tiled, bordered and exploding windows, with size and numbers limited only by RAM.

Menus. Choose from pop-up, pull-down, 123-style, or slug menus, or create your own.

Context-sensitive help. Link help messages to individual screens or fields. Cross reference messages to create hypertext-like help.

Screen design. Build any type of screen or form with the Look and Feel™ Screen Designer, then automatically convert it to C.

Screen flexibility. Call screens from files at run time or link them in.

OAKLAND

Oakland Group, Inc. 675 Massachusetts Ave., Cambridge, MA 02139 USA. FAX: 617-868-4440; **Washington** 206-746-8767; **Benelux** (02159)46814; **Denmark** (02)88 72 49; **France** (1)46 09 28 28; **Germany/Austria/Switzerland** (49)07127/5244; **Norway** (02)44 88 55; **Sweden** (013)124780; **U.K.** (0992)500919. C-scape and Look & Feel are trademarks of Oakland Group, Inc. MS-DOS and XENIX are trademarks of Microsoft Corp. OS/2 is a trademark of International Business Machines Corp. UNIX is a trademark of AT&T. HERCULES is a trademark of Hercules Computer Technology, Inc. Prices and terms subject to change.
Photo by Jessica A. Boyatt. Kanji by Kaji Aso.

Reader Service Number 5

Do you practice Yourdon/DeMarco Structured Analysis?
Do you need to produce high quality DFD's?
Do you need a tool that fits your budget?

Then MacBubbles™ was made for you!

MacBubbles supports the process of structured analysis with:

- Facilities for creating, modifying, leveling and expanding DFD's
- A Data Dictionary that maintains and lists full where-used information
- Automated balancing checks for DFD's and minispecs

MacBubbles improves productivity:

- Easy to learn
- Easy to use
- Responsive

MacBubbles is economical:

Single copy price	$779.99
Multi-copy discount available	
Demo disk	$25.00

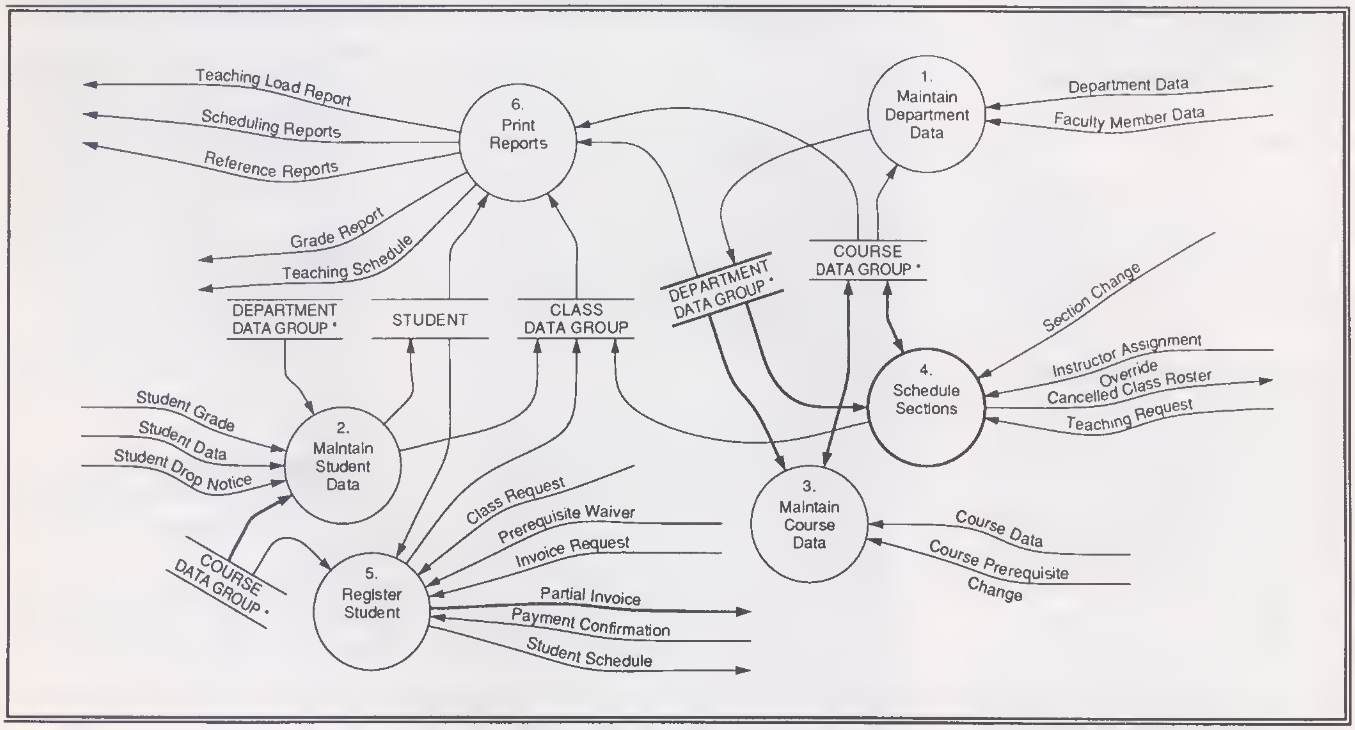

Sample Data Flow Diagram

MacBubbles requires:

- A Macintosh Plus or SE
- Two floppy drives or a hard disk
- A LaserWriter for high quality output

StarSys, Inc.
11113 Norlee Drive
Silver Spring, MD 20902-3619
(301) 946-0522

Reader Service Number 6

Acquisition and Representation of 2D and 3D Data from Turbulent Flows and Flames

Marshall B. Long, Kevin Lyons, and Joseph K. Lam
Yale University

Although turbulence has been studied for over 100 years, prediction of many important quantities in turbulent flows is still not possible, leaving turbulence an unsolved problem of classical mechanics. From an engineering standpoint as well, turbulent flows and flames are extremely important. Most naturally occurring flows are turbulent, as are flames in industrial combustors, automobile engines, jet engines, and rockets. A better understanding of this phenomenon would directly impact a broad range of practical devices and aid in the design of more efficient engines, vehicles with reduced drag, and combustors with lower pollution levels.

Part of the difficulty of the turbulence problem lies in the vast amount of information that must be specified. For example, to fully characterize a turbulent flame would require measurements of temperature, pressure, density, concentrations of all species, and all three components of the velocity vector at each point within a volume. Because of the great diversity of length scales present in a turbulent flow, even a relatively simple turbulent flame could require measurements at more than 10^7 points. At each point, all of the quantities mentioned would have to be recorded as functions of time with characteristic time scales on the order of 10 microseconds.

Laser light-scattering techniques for making two- and three-dimensional measurements, and better methods for representing the large quantities of data obtained may help solve the "turbulence problem."

Clearly, the problem is very complex, and the prospect of determining flow quantities with the measure of detail described above remains distant. Significant progress has been made, however, both theoretically and experimentally. We will present laser light-scattering techniques for making two- and three-dimensional measurements in turbulent flows and flames, along with examples of methods used to represent the large quantities of data obtained. While our work has focused on developing new diagnostic techniques, the methods of data representation and display are topics of immediate importance to scientists hoping to understand the tremendous amount of information their measurements provide. We hope this article will challenge those working in computer processing and graphics by making them more aware of the needs of today's flow and combustion diagnostics community.

Laser diagnostic measurement techniques

In the last decade significant advances in lasers and detector technology, coupled with the advent of laboratory

0018-9162/89/0800-0039$01.00 ©1989 IEEE

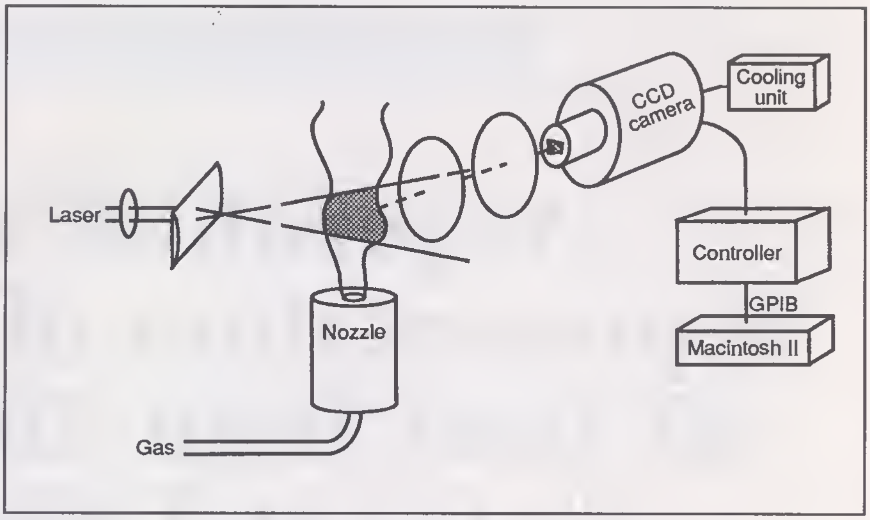

Figure 1. An experimental configuration for two-dimensional light-scattering measurements in turbulent flows and flames. A laser beam is formed into a thin sheet to illuminate the flow. Scattered light is collected and imaged onto a CCD (charge-coupled-device) array detector. The images stored in the computer contain information on temperature, density, or species, depending on the scattering mechanism used.

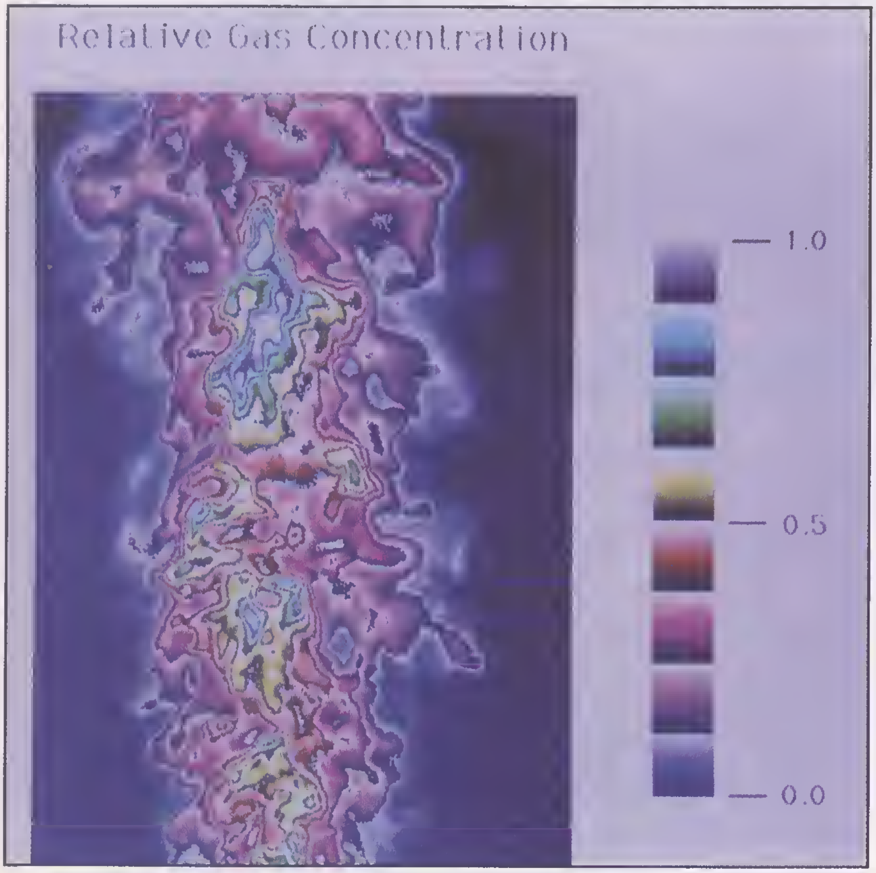

Figure 2. A two-dimensional image of a nonreacting Freon jet mixing with ambient air. The false-color map represents the different Freon concentrations. A discontinuous sequence of hues of varying saturation is used, as the dynamic range of the measurement exceeds the capabilities of easy reproduction by continuous-color variations.

computers to record the data, have prompted development of techniques that have provided a wealth of new data. Rather than relying on disruptive physical probes to make measurements, researchers are making increasing use of the fact that light scattered from the molecules in a flow can reveal a great deal of information. The ability of laser techniques to make nonintrusive, in situ, time-resolved measurements of molecular species, temperatures, velocities, pressures, and densities has been demonstrated.

Figure 1 shows an experimental configuration typical of those used to make two-dimensional measurements in a turbulent flow or flame.[1] The flow used in the examples presented here emerges from a simple round nozzle. Gas issues from the nozzle exit and mixes (and, in some cases, reacts) with the ambient air. To study the turbulent mixing of the gases, a laser beam is formed into a thin sheet that intersects the flow centerline slightly downstream of the nozzle exit. A very small fraction of the laser light is scattered by the flow. This scattered light is collected normal to the illumination sheet by a lens, and an optical filter selects a specific range of wavelengths. The selected wavelength is imaged onto a computer-controlled digital camera, and the intensity distribution is recorded. The particular quantity deduced from the scattered light depends on the specific gases present, the laser power and wavelength, and the detected wavelength. A number of light-scattering mechanisms can be used in this way, including Rayleigh, Lorenz-Mie, fluorescence, and Raman scattering. A recent review article summarized the specifics of these techniques.[2]

Figure 2 shows an example of the data obtained from this type of experiment. In this case the flow is a nonreacting Freon jet mixing with the ambient air. The measurement shows the distribution of Freon within the thin illumination sheet. In the figure a false-color representation is used, with different colors representing different Freon concentrations (that is, mole fractions). Because the dynamic range of the measurement is too great to be easily reproduced by a continuous change in intensity or color, a discontinuous sequence of hues with varying degrees of saturation is used. The color bar shows the mapping of colors to gas concentrations.

As mentioned, many quantities are required to fully characterize a turbulent flame, and current experimental techniques cannot provide them all simultaneously. Some simultaneous measurements can be made, however, by using multiple detectors, with each detector measuring a different quantity by monitoring a different light-scattering mechanism.[3] Figure 3 shows two pairs of concentration and temperature distributions from a turbulent nonpremixed methane-hydrogen flame. The lower part of the figure shows the concentration of methane as determined from Raman scattering, and the upper part shows the corresponding temperature distributions as determined by Rayleigh scattering. The correspondence between the two quantities can be seen by noting that regions of high fuel-gas concentration (shown in yellow on the bottom part of the figure) correspond to areas of lower temperature (blue regions in the upper part of the figure). The information is also complementary, with the temperature distribution giving a better indication of the location of the hot flame zone.

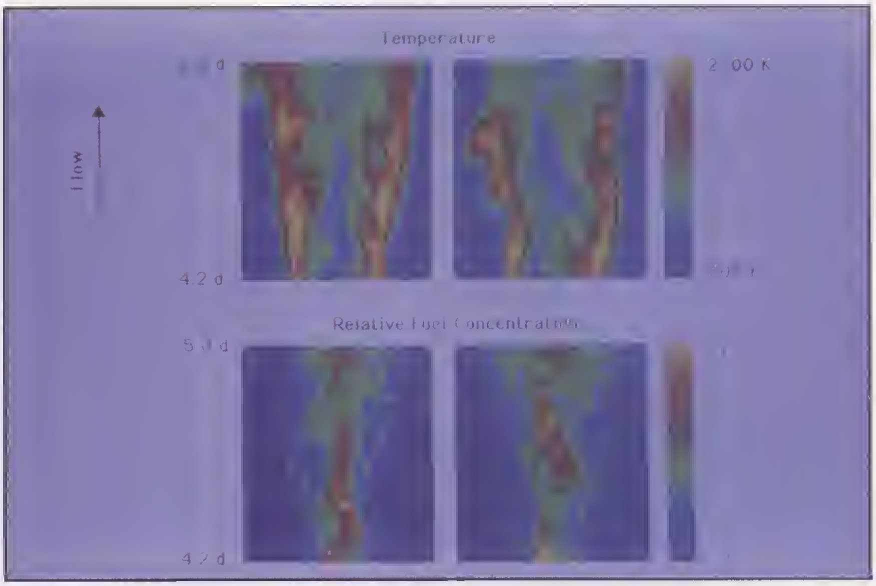

Figure 3. Two pairs of simultaneous temperature and concentration measurements in a turbulent methane-hydrogen flame. The lower part of the figure shows the methane concentration as determined from Raman scattering, while the upper portion shows the corresponding instantaneous temperature distributions as determined from Rayleigh scattering.

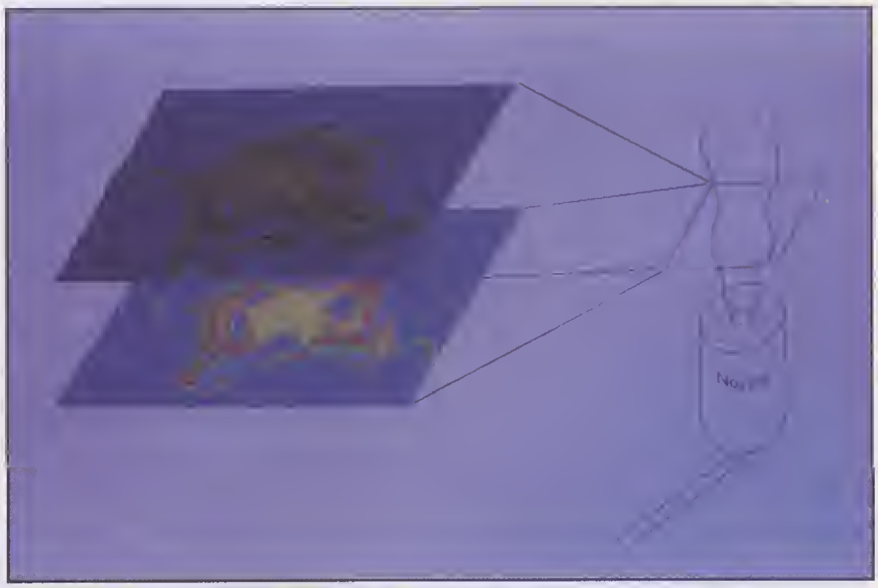

Figure 4. Alternative ways of representing the same two-dimensional gas concentration measurement in a turbulent nonreacting flow. The nozzle gas concentration was measured by imaging Lorenz-Mie scattering in a plane oriented normal to the axis of an axisymmetric jet of Reynolds number 4160. The illumination sheet intersected the centerline six nozzle diameters downstream, where the flow has become three dimensional.

Turbulence by its nature is a three-dimensional phenomenon, so that with data in only two dimensions, some ambiguities remain. For example, Figure 4 shows another visualization of the gas concentration distribution from a turbulent jet. Now, however, the measurement sheet has been oriented normal to the jet

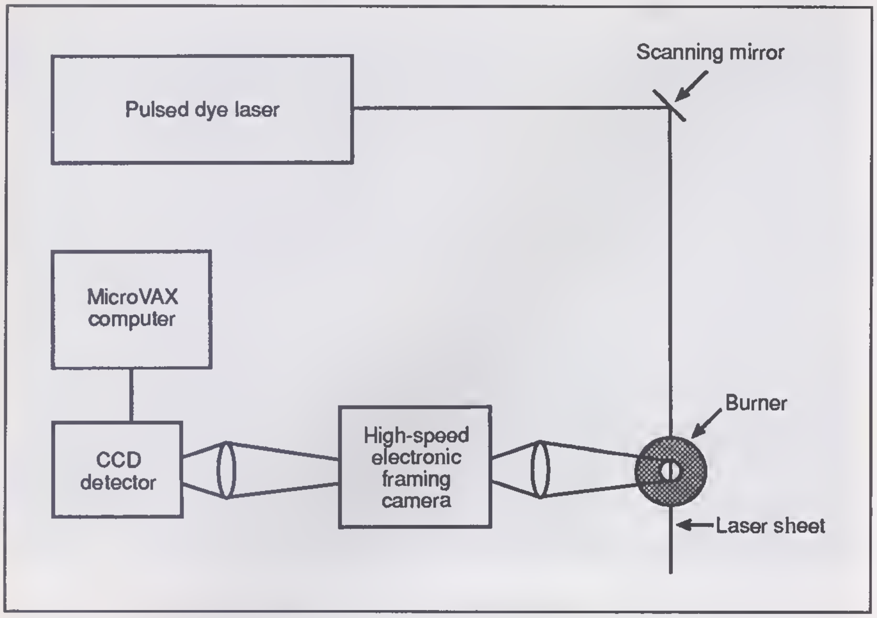

Figure 5. An experimental arrangement used for the time-resolved three-dimensional measurement of a scalar in a gas jet. A laser sheet is swept rapidly through the flow by a rotating mirror. Images of parallel vertical slices are produced on the phosphor screen of a high-speed electronic framing camera. The image is then recorded by a two-dimensional CCD (charge-coupled-device) array. The series of imaged planes constitutes a three-dimensional recording of the flow.

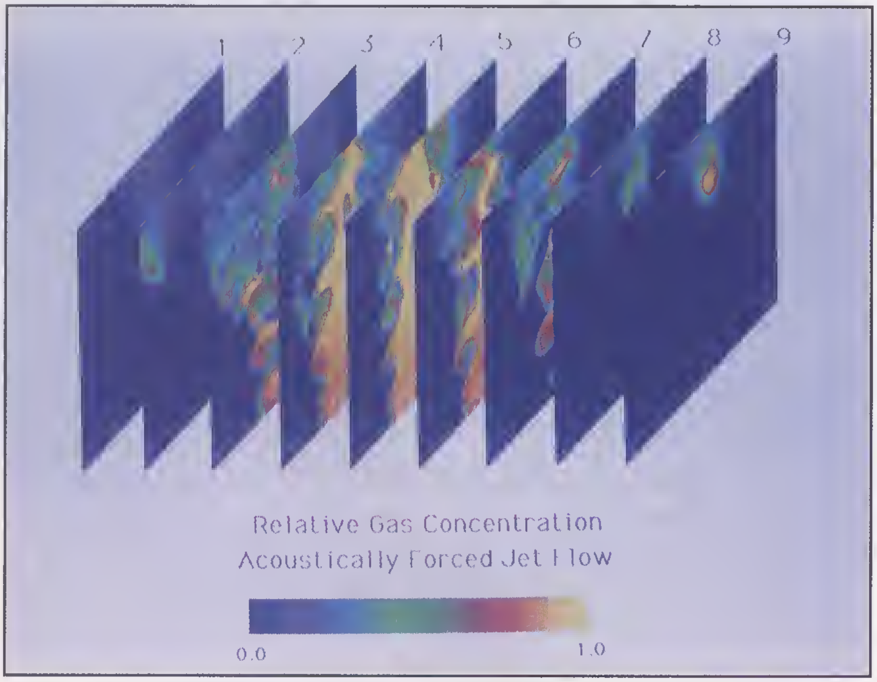

Figure 6. A series of flow planes in an acoustically forced jet, represented by false-color-mapped concentration contours. The concentration is determined from the intensity of Lorenz-Mie scattering. The region imaged is 1.8 centimeters in the streamwise direction and centered 2.4 cm downstream from a nozzle 0.3 cm in diameter. The frames correspond to planes spaced 0.8 millimeter apart.

axis, intersecting it six nozzle diameters downstream. It is clear from Figure 4 that even though the flow emerges from a round nozzle, it quickly loses its axial symmetry, thereby becoming three dimensional. Depending on the orientation of the sheet, very different features of the flow become evident, and thus a single slice through the flow cannot contain all of the information required to fully describe the structures.

Progress has been made in developing techniques capable of providing three-dimensional flow measurements.[4-6] One way to obtain three-dimensional information from a flow is to simply record a large number of two-dimensional measurements from many closely spaced parallel sheets. The most stringent requirement for this measurement is that data from all of the two-dimensional sheets must be recorded before the flow changes significantly. The detection system must have a high framing rate (that is, the rate at which images can be acquired) to record as many two-dimensional sheets as possible during this time. In addition, a laser source is desired that can provide sufficient energy during the measurement time to enable molecular light-scattering processes to be detected.

Figure 5 shows an experimental arrangement that has been used for making essentially instantaneous three-dimensional measurements in nonreacting and reacting flows.[5] A laser illumination sheet is swept rapidly through the flow, and the scattered light from the scanning sheet is recorded by a high-speed electronic framing camera capable of recording 10 to 20 images at framing rates up to 2×10^7 frames per second. The framing camera produces a series of images on a phosphor screen that correspond to different sheet locations. This intensity distribution can then be measured electronically with a CCD (charge-coupled-device) detector. In recent work[6] we showed that it is possible to get truly instantaneous (less than 1 microsecond) gas concentration measurements in turbulent jets and flames.

Figure 6 shows a false-color representation of a series of nozzle gas concentration distributions corresponding to different planes intersecting a flow. The contour characteristics are seen to change as the sheet intersects different locations in the flow. Although the concentration distributions shown in Figure 6 give a complete representation of the three-dimensional data, they do not pro-

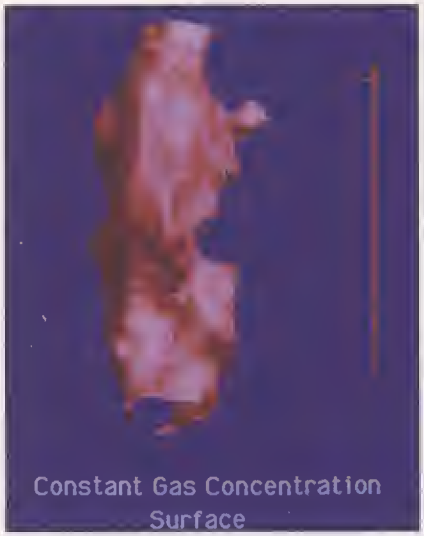

Figure 7. Instantaneous constant fuel-gas concentration surface in a turbulent Bunsen-burner flame. The surface corresponds to a concentration of 64 percent of the maximum in this region and has been generated from experimental data. The flame was a stoichiometric methane-air mixture seeded with biacetyl vapor as a fluorescent marker, with a jet velocity of 19.5 meters per second. The imaged region is 2 × 1 × 0.5 nozzle diameters (*d*=6 millimeters) in size, centered five diameters downstream from the burner. The flow direction is from bottom to top.

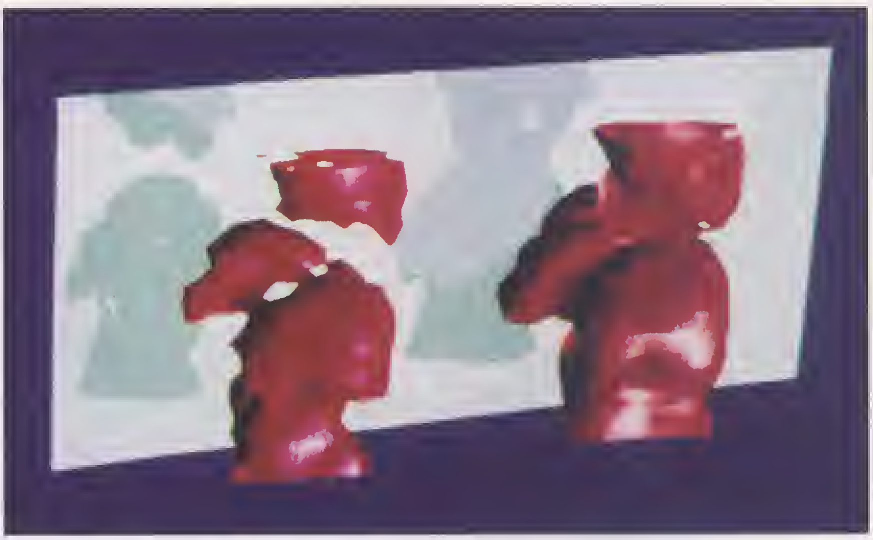

Figure 8. Constant-concentration surfaces in a photoacoustically perturbed Freon-air jet measured by Rayleigh scattering. The left- and right-hand structures correspond to Freon concentrations of 27 percent and 18 percent, respectively. The three-dimensional effect is heightened by the shadows cast.

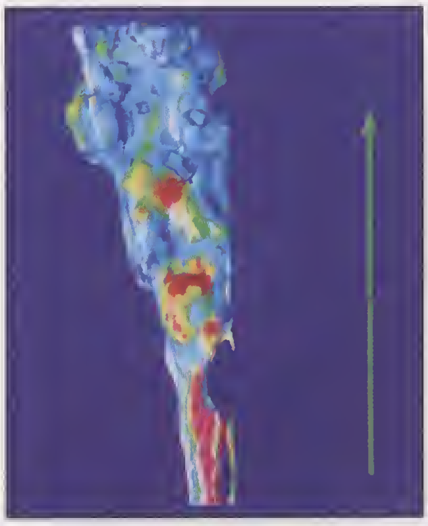

Figure 9. A three-dimensional isoconcentration surface reconstructed from a measurement of Mie scattering from a transitional aerosol-seeded jet of air. The surface corresponds to a region 6.8 × 2.1 × 0.7 nozzle diameters (*d*=2 millimeters) in size, centered 10 diameters downstream from the orifice. The exit velocity of the jet was 8 meters per second. In this representation each point on the surface is color coded according to the magnitude of the concentration gradient vector.

vide much insight into the topology of the flow structures.

The approach we have used is to represent constant-concentration surfaces, which are analogous to lines of constant concentration in a two-dimensional plot. Figure 7 shows a constant gas concentration surface obtained from a set of slices similar to those shown in Figure 6. The data in the figure represent a constant fuel concentration surface in a turbulent premixed flame. Different concentration values would provide different surfaces, with higher concentration contours fitting inside lower ones. Figure 8 shows two different contours obtained from a single three-dimensional data set. Normally, these would lie one inside the other, but, for display, the higher contour value (left) has been removed from inside the lower (right) and light has been cast upon an imaginary plane to heighten the three-dimensional effect.

Figure 9 shows another means of relaying more of the three-dimensional information recorded in a single measurement. The figure shows a constant-concentration surface again, but now the color at each point on the surface has been coded by the magnitude of the concentration gradient vector at that location. The representation of the species data obtained in these measurements as constant-concentration surfaces coincides quite well with one of the current models of turbulent combustion.[7] According to the flame sheet model of combustion, chemical reactions in flames occur in thin sheets located in regions that have the proper mixture of fuel and oxidizer. In this model the data shown in Figure 9 correspond to the flame sheet in a turbulent flame. The availability of three-dimensional data allows the topology of the flame struc-

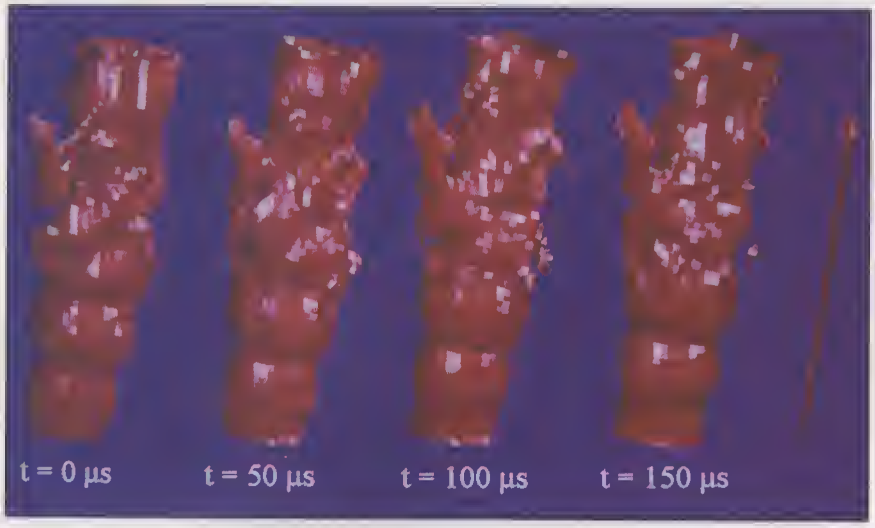

Figure 10. Three-dimensional time evolution of an acoustically forced jet. Shown are instantaneous three-dimensional surfaces obtained by delaying the phase of the measurement relative to the acoustic forcing. The phase is varied at 50-microsecond intervals to produce a ''movie'' sequence. The flow is from bottom to top at a velocity of 8 meters per second. The nozzle diameter is 3 millimeters.

tures to be investigated.

The most difficult aspect of obtaining instantaneous three-dimensional data is the need for the measurement to be made when the flow is essentially stationary. To relax this constraint, measurements can be performed in forced flows. By causing the flow to evolve in a repeatable fashion, the constraint of having to make very rapid measurements is replaced with the need to make the measurement at the right phase of the repeatable flow.[4] Since the flow is repeatable, many instantaneous shots can be accumulated to integrate weak signals. Sequential measurements of several different quantities such as temperature, species, and velocity are also possible.

Another advantage of using forced flows is realized by varying the relative phase of the perturbation and the measurement. In this way the evolution of the three-dimensional structures can be recorded. It is then possible to construct an animated sequence to visualize the development of the three-dimensional surfaces of constant concentration. Figure 10 shows a sequence of surfaces obtained at different phase delays relative to an acoustic perturbation. The convection and evolution of the structures are evident.

Data representation issues

The preceding section gave a sampling of the type of experimental flow data currently available. For most quantitative research in this area, acquisition and display of single images are not sufficient. Ideally, many measurements are made for a given set of flow conditions, and statistical quantities of interest are calculated from the data. Therefore, the amount of data to be stored, analyzed, and displayed is large; effective and efficient means of displaying the data are needed.

The preceding figures showed several display approaches. However, none of these display techniques emerges as the single best approach for visualizing the data. In some cases the physics of the problem may dictate the type of display that most clearly illustrates the phenomenon under study. The data display techniques presented here by no means represent a complete survey of the methods proposed and used to examine flow data.[8] Other possibilities include a flexible scheme that can represent the concentration data as semitransparent clouds.[9] Another attractive approach, though computationally intensive and rather expensive, is to construct a hologram from this measured data.[10]

The constant-property surfaces presented here were generated using a modified version of the commercially available Movie.BYU graphics program. The hidden-surface calculations were performed on a MicroVAX II, and data display was done on a Macintosh II. Because of the complexity of surface rendering and the modest computational power used, calculation of some images took several hours. The surfaces are represented by a polygon mesh, and because of their complexity, many polygons are needed to fully match the resolution of the measurement. With our current hardware-software configuration, the number of polygons that can be represented is limited to approximately 8,000, whereas more than 30,000 polygons would be needed to fully display the resolution of our best measurements.

Another practical issue in displaying the data is cost. Using color provides an important extra dimension in displaying the complex data sets obtained. Currently, however, the cost of publishing two-color plates in some scientific journals represents more than half the cost of the laboratory computer that controls the experiment, stores the data, and displays the results. Therefore, the use of color figures (which can best present the results) might be hard to justify.

Even static color figures are far from the ultimate means of understanding the data. Because of the complexity of some constant-property surfaces, a good object hypothesis is difficult to obtain from a single view. Dynamic projection of three-dimensional data helps considerably. In our laboratory we calculate a sequence of single frames of different object orientations and subsequently animate the precalculated sequence. This animation capability is also extremely useful for time development data like that shown in Figure 10. Ideally, of course, these calculations and the display would be done in real time. Even with the rapid advances in computers, the cost of systems capable of such performance is significant.

Acquiring and displaying data from turbulent flows and flames presents a number of challenges for representing the information we can now obtain experimentally. While the emphasis here has been on experimental data, nearly identical problems face those working in computational fluid dynamics and combustion. Storage, analysis, and display of data represent a significant bottleneck. With the wealth of new data available from both experimental and computational work, there is hope that a more satisfactory means of visualizing data may lead to a significant new understanding of these very complex systems. □

Acknowledgments

We gratefully acknowledge partial support of this research by the Air Force Office of Scientific Research (grant AFOSR-88-0100) and the National Science Foundation (grant MSM 83-51077).

References

1. M.B. Long, B.F. Webber, and R.K. Chang, "Instantaneous Two-Dimensional Concentration Measurements in a Jet Flow by Mie Scattering," *Applied Physics Letters*, Vol. 34, No. 1, 1979, pp. 22-24.

2. R.K. Hanson, "Combustion Diagnostics: Planar Imaging Techniques," *Proc. 21st Int'l Symp. Combustion*, The Combustion Institute, Pittsburgh, 1986, pp. 1,677-1,699.

3. M.B. Long, P.S. Levin, and D.C. Fourguette, "Simultaneous Two-Dimensional Mapping of Species Concentration and Temperature in Turbulent Flames," *Optics Letters*, Vol. 10, No. 6, 1985, pp. 267-269.

4. B. Yip, D.C. Fourguette, and M.B. Long, "Three-Dimensional Gas Concentration and Gradient Measurements in a Photoacoustically Perturbed Jet," *Applied Optics*, Vol. 25, No. 21, 1986, pp. 3,919-3,923.

5. B. Yip et al., "Time Resolved Three-Dimensional Concentration Measurements in a Gas Jet," *Science*, Vol. 235, Mar. 1987, pp. 1,209-1,211.

6. B. Yip, R.L. Schmitt, and M.B. Long, "Instantaneous Three-Dimensional Concentration Measurements in Turbulent Jets and Flames," *Optics Letters*, Vol. 13, No. 2, 1988, pp. 96-98.

7. N. Peters, "Laminar Flamelet Concepts in Turbulent Combustion," *Proc. 21st Int'l Symp. Combustion*, The Combustion Institute, Pittsburgh, 1986, pp. 1,231-1,250.

8. R.A. Drebin, L. Carpenter, and P. Hanrahan, "Volume Rendering," *Computer Graphics* (Proc. SIGGraph 88), Vol. 22, No. 4, Aug. 1988, pp. 65-74.

9. G. Russell and R.B. Miles, "Display and Perception of 3-D Space Filling Data," *Applied Optics*, Vol. 26, No. 6, 1987, pp. 973-982.

10. L. Hesselink, "Digital Image Processing in Flow Visualization," *Ann. Rev. Fluid Mechanics*, Vol. 20, 1988, pp. 421-485.

Marshall B. Long is an associate professor of mechanical engineering at Yale University. His research concerns the development and application of multidimensional laser diagnostics of fluid flow and combustion phenomena. He received a BS in physics from the University of Montana in 1976 and a PhD in applied physics from Yale University in 1980.

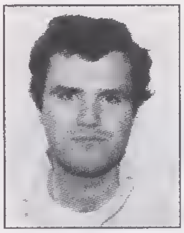

Kevin Lyons is a graduate student in mechanical engineering at Yale University. His research interests include turbulent-flow diagnostics and laser applications. He received a BE in mechanical engineering from Manhattan College in 1988.

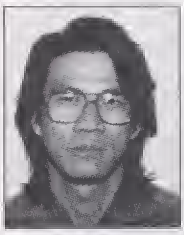

Joseph K. Lam is a research and development technician in the Applied Physics Department at Yale University. His research interests include laser applications and computer imaging. Lam received a BS in physics from Southern Connecticut State University and laser training from San Jose City College.

The authors can be contacted at the Department of Mechanical Engineering, Yale University, PO Box 2157, Yale Station, New Haven, CT 06520.

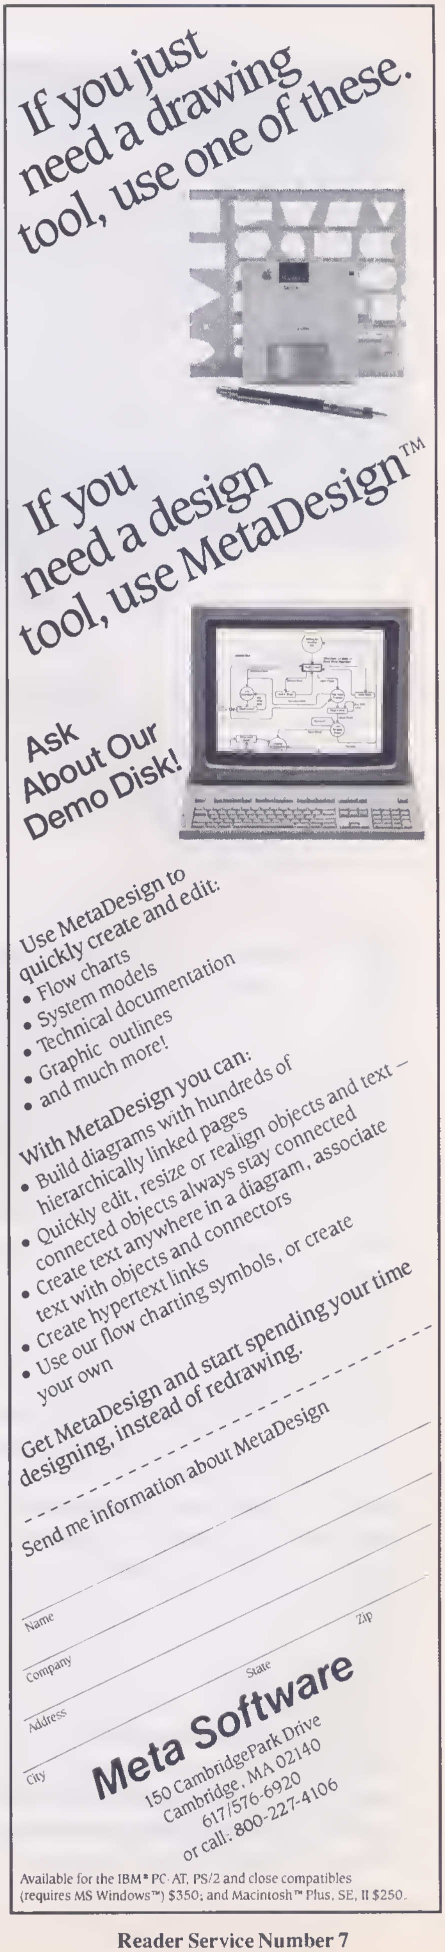

If you just need a drawing tool, use one of these.

If you need a design tool, use MetaDesign™

Ask About Our Demo Disk!

Use MetaDesign to quickly create and edit:
- Flow charts
- System models
- Technical documentation
- Graphic outlines
- and much more!

With MetaDesign you can:
- Build diagrams with hundreds of hierarchically linked pages
- Quickly edit, resize or realign objects and text – connected objects always stay connected
- Create text anywhere in a diagram, associate text with objects and connectors
- Create hypertext links
- Use our flow charting symbols, or create your own

Get MetaDesign and start spending your time designing, instead of redrawing.

Send me information about MetaDesign

Name

Company

Address

City State Zip

Meta Software
150 CambridgePark Drive
Cambridge, MA 02140
617/576-6920
or call: 800-227-4106

Available for the IBM® PC-AT, PS/2 and close compatibles (requires MS Windows™) $350; and Macintosh™ Plus, SE, II $250.

Reader Service Number 7

Interactive Visualization of 3D Medical Data

Henry Fuchs, Marc Levoy, and Stephen M. Pizer
University of North Carolina

Techniques for noninvasively imaging the interior of the human body have undergone a revolution in the last decade. New data acquisition modalities include computed tomography (CT), single photon emission computed tomography (SPECT), positron emission tomography (PET), magnetic resonance imaging (MRI), and ultrasound. All of these modalities have the potential for producing three-dimensional arrays of intensity values.

Unfortunately for the clinician, no fully satisfactory method for viewing this data yet exists. The currently dominant method consists of printing slices of the data onto transparent films for viewing on a backlit screen. This approach makes detection of small or faint structures difficult. It also hampers comprehension of complex three-dimensional structures such as branching arterial systems. Without such an understanding, the planning of surgery, radiation therapy, and other invasive procedures is difficult and error-prone.

Computer-generated imagery offers an effective means for presenting three-dimensional medical data to the clinician. Researchers have, over a period of 15 years, built up a large repertoire of techniques for visualizing multidimensional information. Applying these techniques in clinical settings has until recently been limited, partly because of low image quality and partly because of the large quantity of data involved. The recent advent of larger computer memories and faster processors has spawned a period of rapid growth in software and hardware techniques for data visualization. Widespread clinical use of these techniques can be expected to follow within the next five to 10 years.

Though several techniques for displaying multidimensional data are being developed, research suggests that real-time volume rendering will be preferred for many clinical applications.

Rendering techniques

Three-dimensional medical data has, by itself, no visible manifestation. Implicit in the visualization process is the creation of an intermediate representation—some visible object or phenomenon—that can be rendered to produce an image. This intermediate representation can be almost anything: dots, lines, surfaces, clouds, gels, etc. Since the human perceptual system expects sensory input to arise from physically plausible phenomena and forms interpretations on that basis, the representation should be of a physically plausible object. To promote easier interpretation, it should be of an intuitively familiar one. Most techniques for displaying three-dimensional medical data fall into the three broad categories of surface-based rendering, binary voxel rendering, and volume rendering. The characteristics of each category depend on the intermediate representation employed.

0018-9162/89/0800-0046$01.00 ©1989 IEEE

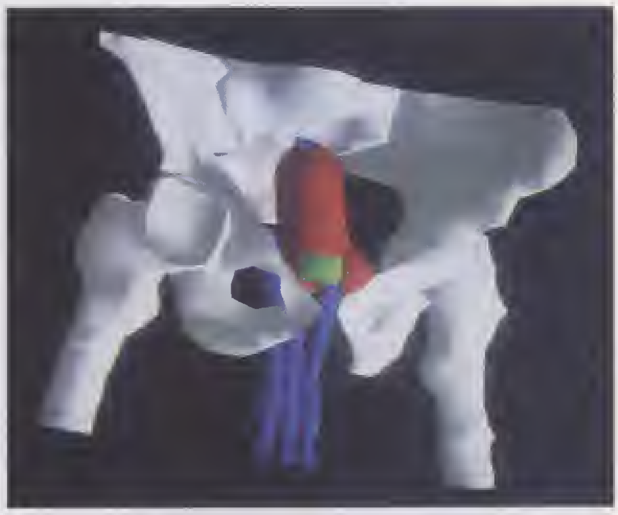

Figure 1. Interactive surface-based rendering on Pixel-Planes 4 of a female pelvis with vaginal inserts for radiation sources. The long bulbous object surrounding the tip of the three vertical shafts is a polygonally defined radiation isodose surface. (CT data courtesy of NC Memorial Hospital)

Surface-based techniques. The first category, surface-based techniques, applies a surface detector to the sample array, then fits geometric primitives to the detected surfaces, and finally renders the resulting geometric representation. The techniques differ from one another mainly in the choice of primitives and the scale at which they are defined. A common technique applies edge tracking on each data slice to yield a set of contours defining features of interest. A mesh of polygons can then be constructed connecting the contours on adjacent slices.[1] Figure 1 was generated using this technique. Alternatively, polygons can be fit to an approximation of the original anatomy within each voxel (a pixel in 3D), yielding a large set of voxel-sized polygons.[2]

These techniques have several desirable characteristics. Geometric primitives are compact, making storage and transmission inexpensive. They also exhibit a high degree of spatial coherence, making rendering efficient. Unfortunately, automatic fitting of geometric primitives to sample data is seldom entirely successful. Contour-following algorithms, for example, occasionally go astray when processing complicated scenes, forcing the user to intervene. These techniques also require binary classification of the incoming data; either a surface passes through the current voxel or it does not. In the presence of small or poorly defined features, error-free binary classification is often impossible. Errors in classification manifest themselves as visual artifacts in the generated image, specifically spurious surfaces (false positives) or erroneous holes in surfaces (false negatives).

Binary voxel techniques. These techniques begin by thresholding the volume data to produce a three-dimensional binary array. The cuberille algorithm renders this array by treating 1's as opaque cubes having six polygonal faces.[3] Alternatively, voxels can be painted directly onto the screen in back-to-front (or front-to-back) order,[4] or rays can be traced from an observer position through the data, stopping when an opaque object is encountered.[5]

Because voxels, unlike geometric primitives, have no defined extent, resampling becomes an important issue. To avoid a "sugar cube" appearance, some sort of interpolation is necessary. These techniques also require binary classification of the sample data, and thus suffer from many of the artifacts that plague surface-based techniques.

Volume-rendering techniques. These techniques are a variant of the binary voxel techniques in which a color and a partial opacity is assigned to each voxel. Images are formed from the resulting colored, semitransparent volume by blending together voxels projecting to the same pixel on the picture plane. Quantization and aliasing artifacts are reduced by avoiding thresholding during data classification and by carefully resampling the data during projection.

Volume rendering offers the important advantage over surface-based or binary voxel techniques of eliminating the need for making a binary classification of the data. This provides a mechanism for displaying small or poorly defined features. Researchers at Pixar of San Rafael, California, first demonstrated volume rendering in 1985. Their technique consists of estimating occupancy fractions for each of a set of materials (air, muscle, fat, bone) that might be present in a voxel, computing from these fractions a color and a partial opacity for each voxel, geometrically transforming each slice of voxels from object space to image space, projecting it onto the image plane, and blending it together with the projection formed by previous slices.[6] Figures 2-4 were generated using an algorithm developed at the University of North Carolina. Though similar in approach to the Pixar technique, the algorithm computes colors and opacities directly from the scalar value of each voxel and renders the resulting volume by tracing viewing rays from an observer position through the data set.[7]

Despite its advantages, volume rendering suffers from a number of problems. High on this list is the technique's computational expense. Since all voxels participate in the generation of each image, rendering time grows linearly with the size of the data set. Published techniques take minutes or even hours to generate a single view of a large data set using currently available workstation technology. To reduce image generation time in our rendering algorithm, we employ several techniques that take advantage of spatial coherence present in the data and its projections. The first technique consists of constructing a pyramid of binary volumes to speed up the search for nonempty (nontransparent) voxels along viewing rays. The second technique uses an opacity threshold to adaptively terminate ray tracing. The third technique consists of casting a sparse grid of rays, less than one per pixel, and adaptively increasing the number of rays in regions of high image complexity. Combining all three optimizations, savings of more than two orders of magnitude over brute-force rendering methods have been obtained for many data sets.[8]

Another drawback of volume rendering is its lack of versatility. Many clinical problems require that sampled data and analytically defined geometry appear in a single visualization. Examples include superimposition of radiation treatment beams over patient anatomy for the oncologist and display of medical prostheses for the orthopedist. We have developed two techniques for rendering mixtures of volume data and polygonally defined objects. The first employs a hybrid ray tracer capable of handling both polygons and sample

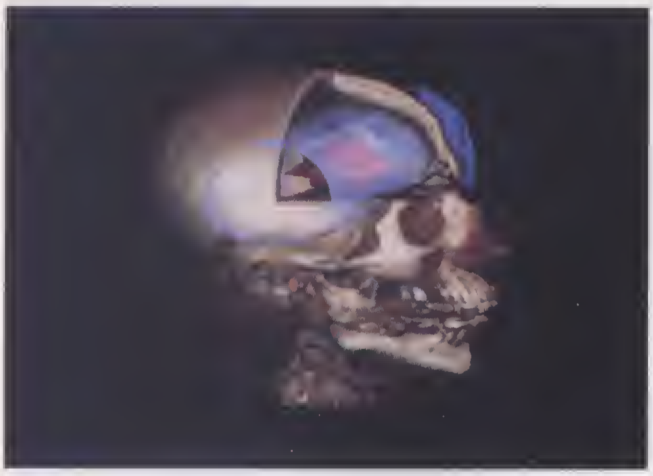

Figure 2. Volume rendering of 256 × 256 × 113-voxel computed tomography study of a human head. A polygonally defined tumor (in purple) and radiation treatment beam (in blue) have been added using our hybrid ray tracer. A portion of the CT data has been clipped away to show the 3D relationships between the various objects. (CT data courtesy of NC Memorial Hospital)

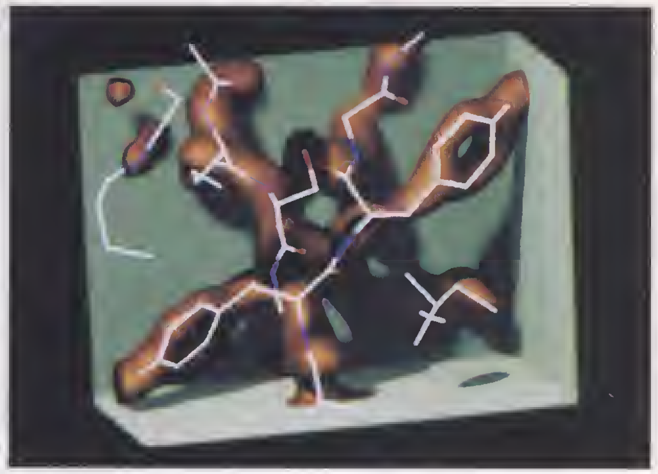

Figure 3. Volume rendering of isovalue contour surface from 24 × 20 × 11-voxel electron density map of staphylococcus aureus ribonuclease. 3D scan-converted backdrop planes, cast shadows, and a color-coded stick representation of the molecular structure have been added to enhance comprehension. Similar techniques can be used to render isodose contour surfaces for radiation treatment planning. (Electron density map courtesy of Chris Hill, University of York)

arrays. Figure 2 was generated using this algorithm. The second consists of a 3D scan conversion of the polygons into the sample array with antialiasing, then rendering the ensemble. Figure 3 includes polygons rendered using this algorithm.[8]

The computation of voxel opacity from voxel value in volume-rendering algorithms performs the dual tasks of classifying the data into objects and selecting a subset of these objects for display. The classification procedure commonly used to display medical data consists of a simple one-for-one mapping from voxel value to opacity. In many cases, it is impossible to obtain the desired visualization using such a mapping. For example, in the magnetic resonance study shown in Figure 4, the scalp and cortical surface have roughly the same voxel value. Any mapping that renders the scalp transparently will also render the cortical surface transparently.

How then to see the cortical surface? One possible strategy is to manually edit the data set to remove the scalp as shown in Figure 5. In this case, the boundary between erased and unerased voxels falls within regions that are rendered transparently. The boundary is therefore not seen in the volume rendering and need not be specified precisely. An alternative solution would be to combine volume rendering with high-level object definition methods in an interactive setting. Initial visualizations made without the benefit of object definition would be used to guide scene analysis and segmentation algorithms, which would in turn be used to isolate regions of interest, producing a better visualization. If the output of such segmentation algorithms included confidence levels or probabilities, they could be mapped to opacity and thus modulate the appearance of the image.

Display hardware

Generating a two-dimensional image from a three-dimensional data set necessarily reduces information content. In particular, images do not use many faculties of the human visual system specifically adapted for perceiving depth: stereopsis, head motion parallax, and the kinetic depth effect. Such specialized display devices as stereo viewers, varifocal mirrors, cine sequences, real-time image-generation systems, and head-mounted displays have been developed to convey as much three-dimensional information as possible to the observer.

Stereo viewers. Relatively simple and inexpensive, stereo viewers offer a means for widening the pathway from the computer to the observer. Our current favorite is a polarizing liquid-crystal plate manufactured by Tektronix. The plate is mounted on the front of a video display, left- and right-eye images are displayed on alternate video fields, and the plate's direction of polarization is electronically switched in synchrony with the alternation of images. The user wears inexpensive passive glasses with polarizing material of different direction in each lens. Using this system, multiple users can see the same stereo image, and each user can look at multiple stereo displays.

Varifocal mirrors. By combining a vibrating reflective surface with a point-plotting CRT, the varifocal mirror provides "true" three-dimensional perception of a distribution of glowing points of light. Our version consists of a vibrating aluminized mylar membrane stretched on a drumhead-like structure and acoustically

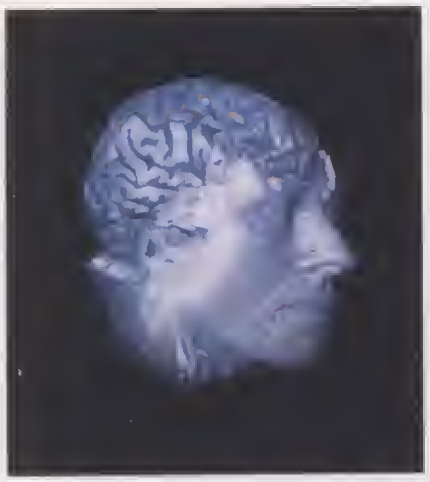

Figure 4. Volume rendering of edited magnetic resonance data set. The apparent mottling of the facial surface in the volume rendering is due to noise in the acquired data. Note that the user was not called upon to define surface geometry, but merely to isolate a region of interest. (MR data courtesy of Siemens AG, edited by Juiqi Tang)

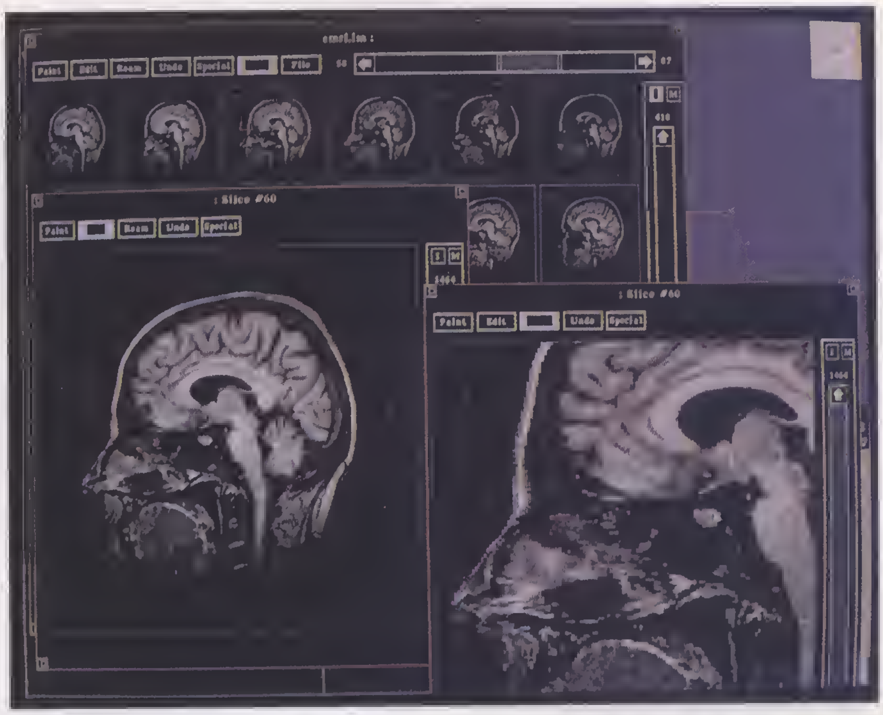

Figure 5. Slices from a 256 × 256 × 156-voxel magnetic resonance study of a human head being edited to remove the scalp and thereby reveal the underlying cortical surface. Depicted is a typical view of the workstation screen in our interactive Imex (image executive) medical-image-editing program, written by Peter Mills.

vibrated from behind by a large conventional speaker driven by a low-frequency (30 Hz) sine wave. A user looking into this vibrating (and "varifocal") mirror sees the image of a simple point-plotting CRT on which a sequence of dots is rapidly displayed. The displayed list of dots is repeated at the 30 Hz rate, synchronized to the mirror vibration. Each dot is thus perceived at a particular depth, depending on the position of the mirror when the dot is displayed. As a user moves his head relative to the display, he sees different views of the data.

This display device provides both stereopsis and head-motion parallax. In addition, it can be viewed simultaneously by several users without their wearing any special apparatus. The limited size of the displayed volume, which is constrained by the size of the CRT and the deflection of the mirror membrane, presents a major drawback. Another disadvantage is the lack of obscuration by the displayed dots; features in the foreground do not hide features in the background.

Cine sequences. Sets of images that have been precalculated, stored in a frame buffer of sufficient memory capacity, and played back in real-time are called cine sequences. The most common type of sequence is one in which the light source and observer are fixed and the object rotates. The resulting play of reflections across surfaces in the data enhances comprehension of the scene. A type of sequence frequently used in volume rendering is one in which the object, the light source, and the observer are all fixed and only the opacities of voxels change. For example, our physician colleagues find it useful to interactively vary the position of a clipping plane, allowing them to study in detail the tiny complex structures of the middle and inner ear. We provide this capability by computing a sequence of images in which the opacities of all voxels lying between the observer and the clipping plane are attenuated to nearly zero, and each image in the sequence has the clipping plane in a slightly different position relative to the data.

When calculating a sequence of images, we usually vary only a single parameter such as observer position, light source position, cutting plane position, or opacity mapping. Users would often prefer to vary multiple parameters. Unfortunately, the number of images that must be precalculated to provide such multiaxis cine sequences is the product of the numbers of steps in the variation of each parameter; a modest 20 steps for each of two variables require the calculation and storage of 400 images. The largest frame buffer currently in our laboratory is the Pixar Image Computer, which has a capacity of 64 512 × 512 12-bit images. This seriously limits the extent to which we can independently vary multiple parameters during playback of cine sequences.

Real-time image-generation systems. Offering more flexibility than cine sequences for the presentation of moving imagery, real-time image-generation systems have found broad acceptance. For surface-based rendering techniques, high-performance graphics workstations are available from Ardent, Silicon Graphics, and Stellar. At the University of North Carolina, we have been using our locally developed Pixel-Planes 4 system to provide clinicians with real-time display of contour-based polygon renderings (see Figure 1).

For binary voxel representations, Phoenix Data Systems, Dimensional Medicine, and Dynamic Digital Displays have fielded systems capable of real-time or near real-time image generation. In addition, Kaufman is currently prototyping a system (and associated algorithms for 3D scan conversion of surfaces and polyhedra) that will be capable of displaying a 512 × 512 × 512 voxel binary volume in real time.[9]

At this writing, no system exists that can perform volume rendering in real time. The most widely used system is the Pixar Image Computer. Using a four-channel, single-instruction, multiple-data processor and Pixar's ChapVolumes software package, the Image Computer can generate high-quality images of 256 × 256 × 256-voxel data sets in about one minute. At the University of North Carolina, we are currently developing Pixel-Planes 5, a massively parallel raster-graphics engine consisting of 16 40-Mflop graphics processors, 250,000 pixel processors, a 512 × 512 pixel color frame buffer, and a 640 megabits-per-second ring network.[10] In addition to enhancing our surface-based rendering capabilities, we intend to integrate Pixel-Planes 5 into a real-time volume-rendering workstation. The frame rate we expect from this workstation depends on what parameters change from frame to frame. Preliminary estimates suggest that for changes solely in observer position, we will be able to generate a sequence of slightly coarse images at between 10 and 20 frames per second. A sequence of images comparable in quality to Figure 4 will be generated at between one and two frames per second.

Head-mounted displays. Pioneered by Ivan Sutherland at the University of Utah in 1968, head-mounted displays incorporate stereopsis, head-motion parallax, and realistic images (including obscuration). These devices present stereo images to the user through displays mounted in front of each eye, track the motion of the user's head, and generate new images in real time in response to changes in the user's head position and direction of view. This provides the illusion of walking about a simulated object or even of walking inside a simulated environment. Most of the development in this area has been for military and aerospace applications at such places as Wright Patterson Air Force Base and NASA Ames Research Center. We have also been developing head-mounted displays for several years.

These systems demand very fast update rates and currently suffer from several technological weaknesses: poor resolution of the small video displays, high latency in the tracker-to-image feedback loop, and limited-range tracking systems. Nonetheless, they continue to hold great promise, and we hope to eventually incorporate a head-mounted display, Pixel-Planes 5, and our volume-rendering algorithms into an interactive radiotherapy treatment-planning system.

Topics for future research

A vital component of any computer-assisted medical image display system is a satisfactory user interface. In addition to the firms and research groups already mentioned, interactive medical image display systems have been developed by CEMAX, the Mayo Clinic, Sun Microsystems, and Lee Westover of our own department. User interfaces for many existing systems are constrained by the inability to generate images in real time. For volume-rendering systems in particular, feedback during selection of rendering parameters is usually provided by metavisualizations such as 2D plots of color and opacity versus input value and wire-frame representations of viewing frustums and motion paths. If future systems succeed in generating volume-rendered images in real time, these metavisualizations can be omitted or relegated to a supporting role. In such systems, sequences of rendered images would serve as feedback to the user of changes made in rendering parameters.

An important problem that has only begun to be addressed is how to visualize multiple data sets. Examples include superimposition of anatomic information from CT data and metabolic information from PET data, or simultaneous visualization of anatomy and radiation dose for cancer therapy planning. Realistic shading models have been used to visualize multiple 2D scalar fields by assigning one data set to surface relief and another to surface albedo or color. For three-dimensional data, the second data set can be used to modulate the color or opacity computed for voxels in the first data set. Alternatively, the second data set could be used to perturb the normals, shininess, or other properties of the surfaces displayed from the first data set. Recent investigations in this area by Hoehne et al.[11] and workers at the University of Chicago[12] show promise.

Current perspective suggests that the applicability of volume rendering to certain disciplines—diagnostic radiology in particular—hinges on answering the question: How correct is a volume rendering? Volume rendering, relative to surface-based techniques, has the disadvantage of rendering fuzzy surfaces represented by a band-limited sample array rather than sharply defined surfaces represented by geometric primitives. Since fuzzy surfaces do not occur in daily life, they are not intuitively familiar. Moreover, volume-rendered surfaces do not look fuzzy; they look precise. The danger therefore exists of interpreting these visualizations incorrectly. Since the radiologist is often required to make subtle judgments of feature size and shape, the potential exists for making erroneous diagnoses based on these visualizations. Clearly, volume rendering is useful for navigating through complex 3D data and for identifying gross anatomic abnormalities. It remains to be seen how useful it can be for perceiving fine features and to what extent it can replace slice-by-slice presentation of medical data in clinical practice.□

Acknowledgments

We thank our collaborators in the Radiation Oncology Department at UNC School of Medicine and NC Memorial Hospital, especially Julian Roseman, Ed Chaney, and George Sherouse. We also thank Andrew Skinner and Richard David for assistance with clinical studies, Phil Stancil for hardware systems support, and Rick Pillsbury for observations about the utility of cine sequences. This research is partially supported by NIH Grant No. 1 P01 CA47982-01, NSF Grant No. CDR-86-22201, DARPA ISTO Order No. 6090, ONR Contract No. N00014-86-K-0680, and IBM.

References

1. H. Fuchs, Z.M. Kedem, and S.P. Uselton, "Optimal Surface Reconstruction from Planar Contours," *Comm. ACM*, Vol. 20, No. 10, Oct. 1977, pp. 693-702.
2. W.E. Lorensen and H.E. Cline, "Marching Cubes: A High-Resolution 3D Surface Construction Algorithm," *Computer Graphics* (Proc. SIGGraph), Vol. 21, No. 4, July 1987, pp. 163-169.
3. G.T. Herman and H.K. Liu, "Three-Dimensional Display of Human Organs from Computer Tomograms," *Computer Graphics and Image Processing*, Vol. 9, No. 1, Jan. 1979, pp. 1-21.
4. G. Frieder, D. Gordon, and R.A. Reynolds, "Back-to-Front Display of Voxel-Based Objects," *IEEE Computer Graphics and Applications*, Vol. 5, No. 1, Jan. 1985, pp. 52-59.

5. D.S. Schlusselberg and W.K. Smith, "Three-Dimensional Display of Medical Image Volumes," *Proc. NCGA 86*, Nat'l Computer Graphics Assn., Fairfax, Va., Vol. III, May 1986, pp. 114-123.
6. R.A. Drebin, L. Carpenter, and P. Hanrahan, "Volume Rendering," *Computer Graphics* (Proc. SIGGraph), Vol. 22, No. 4, Aug. 1988, pp. 65-74.
7. M. Levoy, "Display of Surfaces from Volume Data," *IEEE Computer Graphics and Applications*, Vol. 8, No. 3, May 1988, pp. 29-37.
8. M. Levoy, *Display of Surfaces from Volume Data*, doctoral dissertation, University of North Carolina at Chapel Hill, May 1989.
9. A. Kaufman and R. Bakalash, "Memory and Processing Architecture for 3D Voxel-Based Imagery," *IEEE Computer Graphics and Applications*, Vol. 8, No. 6, Nov. 1988, pp. 10-23.
10. H. Fuchs et al., "Pixel-Planes 5: A Heterogeneous Multiprocessor Graphics System Using Processor-Enhanced Memories," *Computer Graphics* (Proc. SIGGraph), Vol. 23, No. 4, July 1989.
11. K.H. Hoehne et al., "Display of Multiple 3D-Objects Using the Generalized Voxel-Model," *Proc. SPIE*, Vol. 914, 1988, pp. 850-854.
12. D.N. Levin, et al. "Integrated 3D Display of MR, CT, and PET Images of the Brain," *Proc. NCGA 89*, Nat'l Computer Graphics Assn., Fairfax, Va., Vol. 1, pp. 179-186.

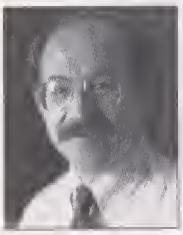

Henry Fuchs is the Federico Gil Professor of Computer Science and an adjunct professor of radiation oncology at the University of North Carolina at Chapel Hill. His research interests include interactive 3D graphics, VLSI architectures, and 3D medical imaging. The principal investigator for several research projects funded by the Air Force, DARPA/ISTO, NIH, and NSF, he also served as chairman of the 1985 Chapel Hill Conference on VLSI and the 1986 Workshop on Interactive 3D Graphics held at UNC-Chapel Hill. He is on several advisory committees, including NSF Division of Microelectronics Information Processing Systems and Stellar Computer's Technical Board. He received a BA from the University of California at Santa Cruz in 1970 and a PhD from the University of Utah in 1975.

All three authors can be contacted at the Department of Computer Science, Sitterson Hall, University of North Carolina, Chapel Hill, NC 27599-3175

Marc Levoy is a research assistant professor of computer science and radiation oncology at the University of North Carolina at Chapel Hill. He is the principal developer of the Hanna-Barbera Computer Animation System and served as its director from 1980 through 1982. His research interests include scientific visualization, texture mapping, volume rendering, medical imaging, and molecular graphics. He received his bachelor's and master's degrees from Cornell University in 1976 and 1978 and his PhD in computer science from the University of North Carolina at Chapel Hill in 1989.

Stephen M. Pizer is a professor of computer science, radiation oncology, and radiology at the University of North Carolina at Chapel Hill, where he heads the multidepartmental Medical Image Display Research Group. His research covers human and computer vision, interactive 3D graphics, contrast enhancement, and workstation development. He conducts active collaborations with laboratories in the Netherlands, Germany, and the US, and is associate editor for display of *IEEE Transactions on Medical Imaging*. He received a PhD in computer science from Harvard in 1968.

**Universität
Gesamthochschule Paderborn**

On the basis of agreements between the Westphalia Foundation, the State of Northrhine-Westphalia and the University of Paderborn,

The Heinz-Nixdorf-Institute

Interdisciplinary Research Centre for Information Technology (CIT) is being created within the University.

The task of CIT is to strengthen teaching and research at the university in application-oriented areas of informatics and technology including scientific foundations and topics from the fields of the arts and social sciences. CIT is a central academic institute of the University. The position is a full professorship integrated into the Department involved and included in the curriculum of relevant courses of study. It is the aim of CIT to stimulate innovation in research by means of joint projects between researchers from different fields. The affiliation of a centre for postgraduate research is one of the components of the agreement between the Westphalia Foundation and the State of Northrhine-Westphalia.

The selection of applicants will be undertaken by the Department in cooperation with the Westphalia Foundation and the Board of Founding Directors of CIT. Candidates are expected to show particular willingness to cooperate in joint projects between different fields.

The foundation period as planned requires the following position to be filled at the next possible date:

Department 17 - Mathematics/Informatics

A University Chair (Grade C4) in Practical Informatics,
particularly Practical Studies in Distributed Systems

Applications should be qualified in one of these fields: "Computer Networks/Operating Systems", "Software Engineering for Distributed Systems" or "Distributed Databases". Within the framework of the research topics of CIT, the successful candidate will be expected to pay particular attention to problems involved in the development of distributed systems. Practical experience in the implementation of relevant projects is an advantage.

Required qualifications:
Section 49 para. 1 no. 4a, WissHG NW (University Law of Northrhine-Westphalia): Professorial thesis or equivalent academic/scientific qualification inside or outside university.

Application deadline: 6 weeks after publication.

The University of Paderborn is seeking to increase the proportion of women in university teaching positions and expressly requests women to apply.

Severely handicapped applications with equivalent suitability to others will be given priority.

Applications are invited to be presented with the usual documents (curriculum vitae, certificates, list of publications in major scientific periodicals) stating the code number 403 to the dean of the Department 17 - Mathematics/Informatics, University of Paderborn, Postfach 1621, D-4790 Paderborn.

Developers from Boston to Yokohama choose Design/OA™ to build graphic applications. *Fast.*

Chiyoda Inc. Yokohama, Japan
Sapphire Software San Francisco, CA
Navigator Software Santa Barbara, CA
Naval Ocean Systems Center San Diego, CA
Evans & Sutherland Salt Lake City, UT
University of Kansas Lawrence, KS
University of Michigan Ann Arbor, MI
US Air Force Wright-Patterson AFB
SWIFT, Inc. Culpeper, VA
RPI Troy, NY
University of Delaware Newark, DE
Mitre Corporation McLean, VA
SRS Technologies McLean, VA
PRC McLean, VA
MIT Cambridge, MA
Honeywell-Bull Waltham, MA
McCormack & Dodge Natick, MA
Community Publication Network Hull, MA
Marine Midland Bank New York, NY
SUNY Stony Brook, NY
AT&T Bell Labs Middletown NJ
University of Edinburgh Edinburgh, Scotland
IGL Technology Paris
Universite de Paris-Sud Orsay, France
University of Zaragoza Zaragoza, Spain
LAAS du CNRS Toulouse, France
Devtech International Heerlen, The Netherlands
GMD Bonn
Pronik AG Biel/Bienne, Switzerland
Aarhus University Aarhus, Denmark
Deutsche Bundespost Darmstadt, W. Germany
Universita De Milano Milano, Italy
Telenokia Oy Espoo, Finland
Helsinki University Espoo, Finland
Abo Akademi Abo, Finland
Fraunhofer-Institut Berlin
Institut fur Informatick Munich
Datamat Roma, Italy

Honeywell-Bull, Waltham MA

Honeywell-Bull used Design/OA to create a graphical micro-computer interface to monitor their nightly mainframe batch processing. The application retrieves batch processing data, initiates the batch jobs, and provides the operator real-time processing status throughout the night.

PRC, McLean VA

Planning Research Corporation used Design/OA to develop an Ada diagramming tool by adding Ada specific graphics to Design/2.0. This new tool, called AdaDesign, is used by Ada programmers for flow charting and system design.

MIT, Cambridge, MA

"With Design/OA we can go from an idea to an implementation very quickly. It has expanded the range of questions we can explore, given the same resources."
Alexander Levis, Senior Research Scientist, Laboratory for Information and Decision Systems.

AT&T Bell Labs, Middletown NJ

AT&T Information Systems used Design/OA to build a reverse engineering application that provides dynamic documentation of large mainframe C programs. Programmers can graphically view the program's functions, files, variables, and their relationships.

Design/OA, The Open Architecture Development System from Meta Software, is the graphics development shell that enables you to create graphic applications in a fraction of the time required by other methods.

Design/OA is used to build tools for software engineering, system design, structured analysis, and system simulation. Virtually any system design methodology can be represented graphically with Design/OA. Design/OA is also an effective means of creating graphic interfaces to other micro, mini or mainframe applications and databases.

Complex graphic relationships can be rapidly defined and manipulated, both in overview and across lower levels of detail. You need only work with high-level support tools, enhancing development productivity and minimizing implementation time, without changing your development style. These tools also make it easy to maintain or enhance your application.

Design/OA is available for MS-DOS®/MS-Windows,™ Macintosh™ OS, and in the Fall of 1989, UNIX™/X Window System.™ Applications built with Design/OA are portable across all three platforms. Meta Software offers Design/OA users both initial training support and advanced implementation services.

To explore the potential of Design/OA for your application, write Andrew Levin, Meta Software Corp., 150 CambridgePark Drive, Cambridge, MA 02140, or call 800-227-4106 (617-576-6920 in MA).

Meta Software

© 1988 Meta Software Corp. Design/OA is a trademark of Meta Software Corp. MS-DOS is a registered trademark and MS-Windows is a trademark of Microsoft Corporation. Macintosh is a trademark of Apple Computer, Inc. UNIX is a trademark of AT&T. X Window System is a trademark of MIT.

Reader Service Number 9

Visualizing Large Data Sets in the Earth Sciences

William Hibbard and David Santek
University of Wisconsin-Madison

With single data sets containing billions of points, meteorologists and other earth scientists face an avalanche of data received from their remote-sensing instruments and numerical simulation models.

The Space Science and Engineering Center (SSEC) at the University of Wisconsin-Madison led the evolution of weather visualization systems from paper to electronic displays with our Man-computer Interactive Data Access System (McIDAS). We have continued this evolution into animated three-dimensional images and recently into highly interactive displays. Our software can manage, analyze, and visualize large data sets that span many physical variables (such as temperature, pressure, humidity, and wind speed), as well as time and three spatial dimensions.[1,2]

McIDAS has produced three-dimensional animations of data sets from many diverse sources on videotape and on a special binocular stereo workstation. These animations are used in the classroom at the University of Wisconsin and other institutions, and by visiting scientists presenting case studies at conferences. Model developers are using McIDAS to produce highly interactive, real-time animations of the dynamics of their models.

McIDAS, an interactive visualization system, is revolutionizing the ability of earth scientists to manage and analyze data from remote sensing instruments and numerical simulation models.

Data management

The McIDAS system manages data from at least 100 different sources, a diversity illustrating the flexibility of our data management tools. These tools consist of data structures for storing different data types in files, libraries of routines for accessing these data structures, system commands for performing housekeeping functions on the data files, and reformatting programs for converting external data to our data structures.

McIDAS includes data structures for grids, images, paths, and nonuniform data. Two- and three-dimensional grids are spatial arrays of numbers appropriate to numerical weather-model output and analyses of remote sensed data. Our grid structures permit grouping a range of times and multiple physical variables in a data set of billions of points, allow a variety of horizontal and vertical map projections, and indicate missing data points. Images are two-dimensional arrays of pixel intensities produced by satellites, radars, lidars (laser radars), and other sources. These images may include multiple spectral bands and may be grouped into time sequences, possibly containing billions of points. Paths are sequences of points through space or space-time, usually representing trajectories of air parcels derived from wind data. Nonuniform data consist of numbers without any spatial order that might be produced by surface observations, balloons, ships, and aircraft.

The McIDAS system provides numerous analysis functions for time and space interpolation of data, for deriving grids from nonuniform data, for deriving trajectories from wind grids, for converting data to a moving frame of reference, for making measurements from image data, and for applying a large variety of operators to grid and image data. (As part of the development of our three-dimensional visualization tools, we have inventoried four-dimensional data sets in meteorology and related earth sciences. A copy of this inventory is available at cost by writing to the librarian of SSEC.)

0018-9162/89/0800-0053$01.00 ©1989 IEEE

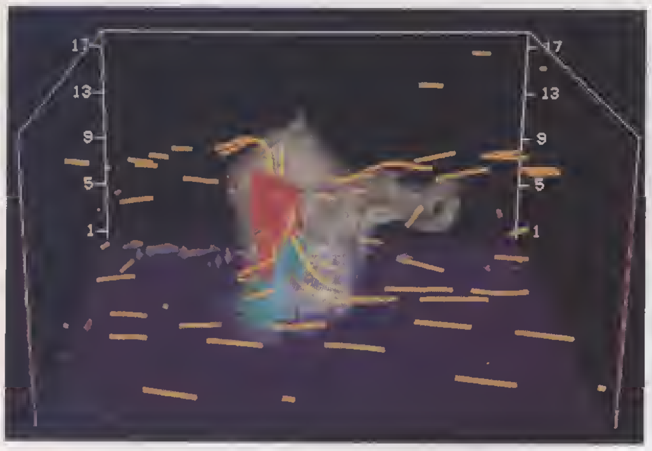

Figure 1. Regional Atmospheric Modeling System (RAMS) cloud model of a severe thunderstorm, showing a transparent 1.0-gram-per-cubic-meter condensate surface, a red 4.0-gram-per-cubic-meter hail surface, a blue 2.0-gram-per-cubic-meter rain surface, and wind trajectories. (With Gregory Tripoli, Univ. of Wisconsin Meteorology Dept.)

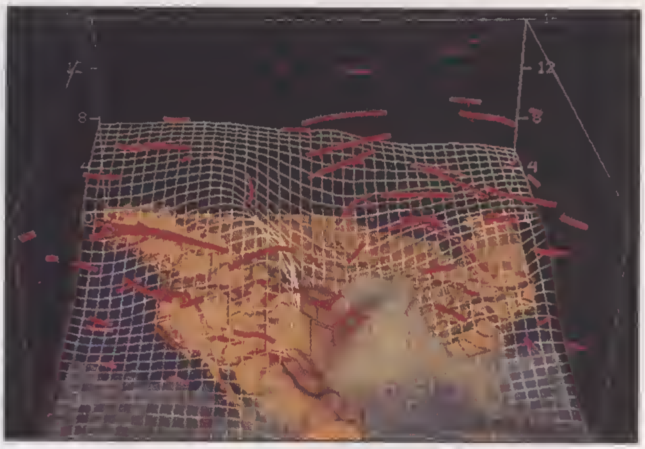

Figure 2. Limited-Area Mesoscale Prediction System (LAMPS) model output showing a grid-mesh 300-kelvin potential temperature surface, a transparent 8-gram-per-kilogram mixing-ratio moisture surface, and wind trajectories over a topographical map. (With Patricia Pauley, Univ. of Wisconsin Meteorology Dept.)

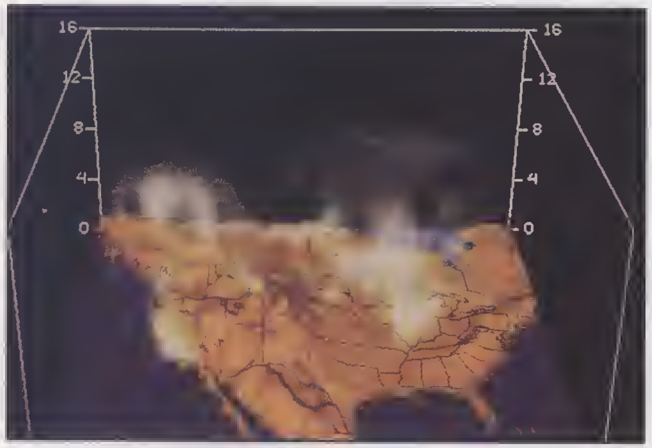

Figure 3. LAMPS model output for cloud water density over a topographical map. Volume is rendered as a transparent fog, with opacity proportional to cloud water density. (With Patricia Pauley, Univ. of Wisconsin Meteorology Dept.)

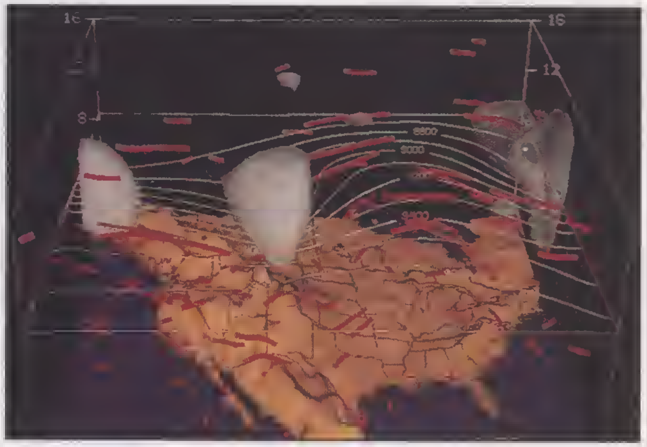

Figure 4. LAMPS model output showing a transparent 0.00016-per-second vorticity surface, trajectories, and 300-millibar height contour lines over a topographical map. (With Patricia Pauley, Univ. of Wisconsin Meteorology Dept.)

Three-dimensional visualization

The McIDAS tools for three-dimensional visualization of meteorological data run on an IBM mainframe and can load up to 128-frame animation sequences into workstations built at SSEC. The animations can be viewed directly (optionally, in binocular stereo) or recorded on videotape.

Our data are depicted in a rectangular box outline with height tick marks labeled in kilometers. A map, including topographical relief and boundary lines, can be drawn on the bottom of the box. The box visually defines the three-dimensional space, the map provides geographical context, and the topography is useful because it affects the weather.

We can depict gridded three-dimensional scalars using opaque or transparent contour surfaces (Figure 1) or a grid-mesh contour surface (Figure 2). A contour surface defines a constant value for a scalar variable. McIDAS can also depict a scalar as a transparent fog (Figure 3), where fog opacity is proportional to any scalar variable. The system generates contour lines (Figure 4) on a two-dimensional slice through a three-dimensional gridded scalar.

Wind vectors can be depicted using trajectory paths, which can be long and tapered (Figure 5), short with length proportional to speed (Figure 6), or faded to transparency (Figure 7). McIDAS also represents winds with derived scalars, such as vorticity, contoured in Figure 4, or potential vorticity, contoured in Figure 5. Figure 8 shows depth-cued trajectory paths adapted to depict a network of underground caves.

McIDAS can texture map image data to surfaces in three-dimensional perspective, as shown in the two lidar images in Figure 9. In Figure 10, we have texture-mapped a visible satellite image to a surface defined by the corresponding infrared image. Pattern recognition discriminates clouds, and the visible pixels are blended with an artificial shade based on cloud-top geometry.

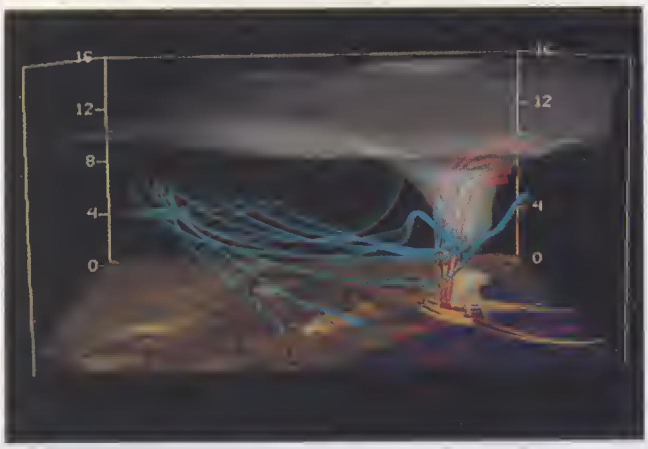

Figure 5. Mesoscale Analysis and Simulation System (MASS) model simulation of the Presidents' Day storm, showing a transparent 0.00002-kelvin per millibar per second potential vorticity surface and trajectories over a topographical map. (With Louis Uccellini, NASA Goddard Space Flight Center)

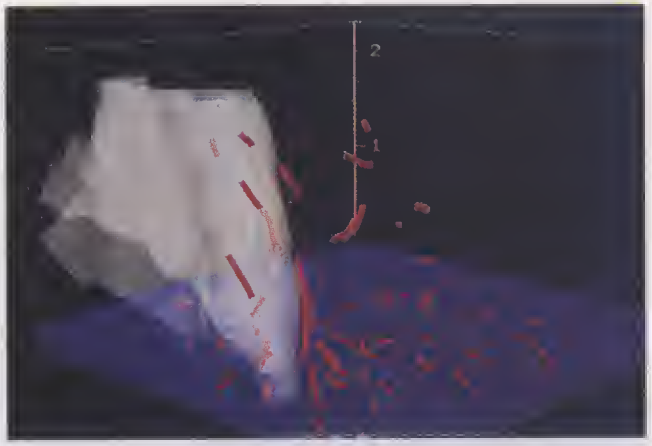

Figure 6. A microdownburst sensed by multiple doppler radars, with a transparent 30-DBZ radar reflectivity surface and wind trajectories. (With Robert Kropfli, NOAA Environmental Research Lab.)

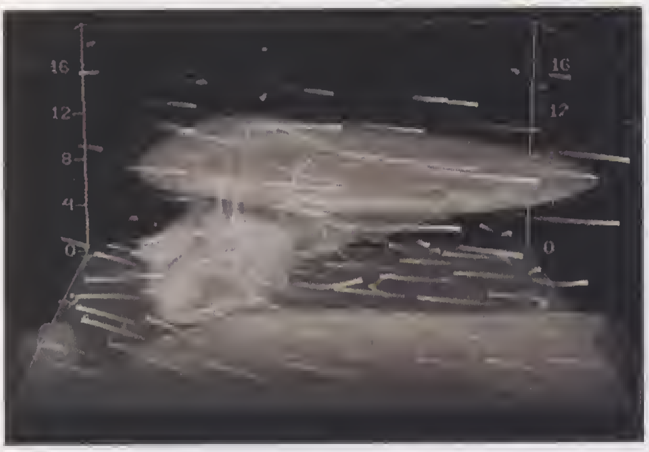

Figure 7. Thunderstorm simulation showing a transparent 0.5-gram-per-cubic-meter cloud water surface and trajectories. (With Robert Schlesinger, Univ. of Wisconsin Meteorology Dept.)

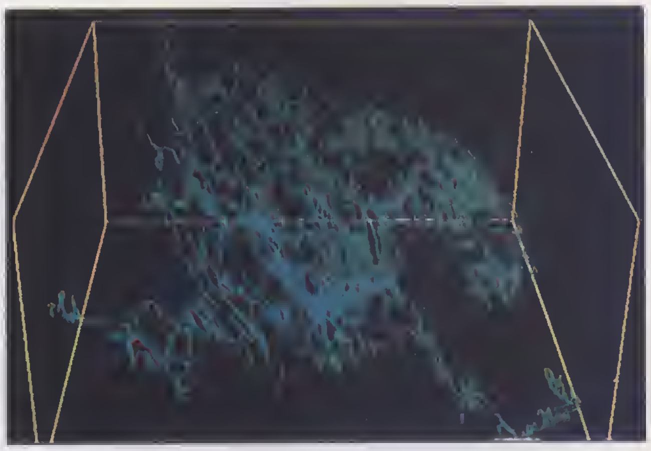

Figure 8. The Wind Cave network from data gathered by manual survey. (With Jem Nepstad, US Nat'l Park Service)

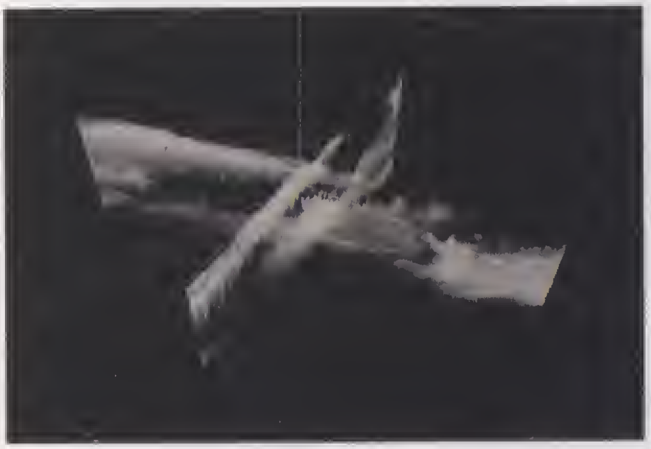

Figure 9. Cirrus cloud section by two vertical lidar (laser radar) slices. With Ed Elorant, Univ. of Wisconsin Meteorology Dept.)

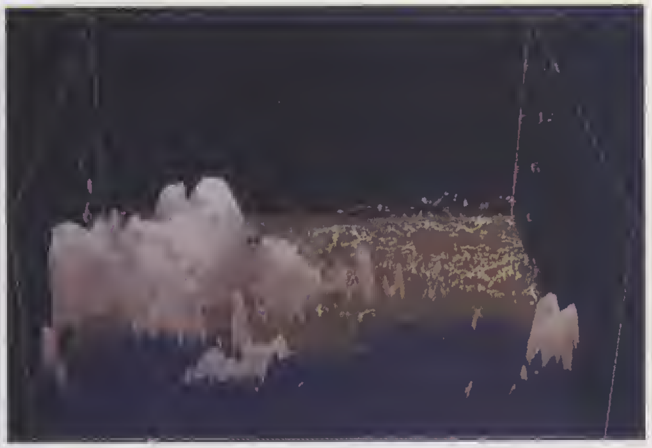

Figure 10. Perspective image of clouds over the Gulf of Mexico, generated from Geostationary Operational Environmental Satellite (GOES) infrared and visible data.

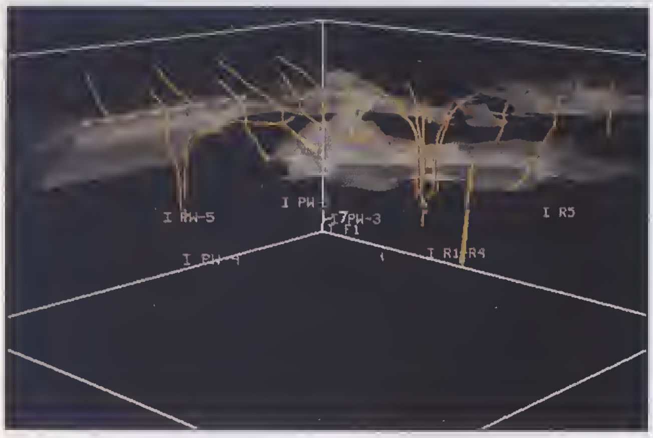

Figure 11. Hydrological model output showing transparent clay layers, ground-water trajectories, and labeled well positions. (With Kenneth Quinn, Warzyn Engineering)

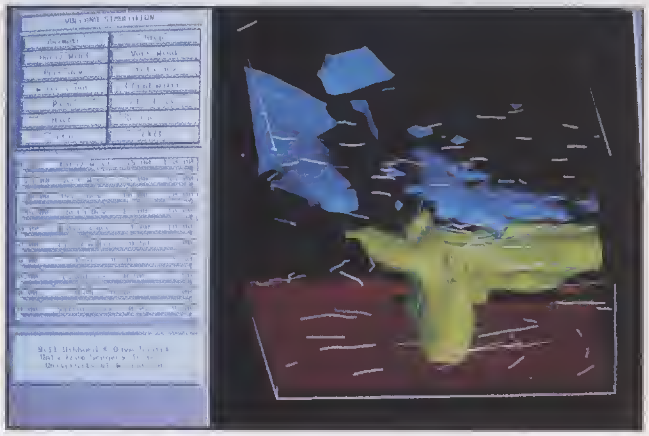

Figure 12. RAMS model of the atmospheric effects of a volcanic eruption, showing a green 0.07-gram-per-cubic-meter sulfur dioxide surface, a red 4.0-gram-per-cubic-meter water vapor surface, a blue 12.0-kelvin potential temperature deviation surface, and wind trajectories. The widgets to the left of the image control the display. (With Gregory Tripoli, Univ. of Wisconsin Meteorology Dept.)

Our images also show the interactions between physical quantities. Figure 1 shows contour surfaces for condensate (transparent), hail (red), rain (blue), and wind trajectories from a thunderstorm simulation. Figure 2, taken from a model analysis of an extratropical cyclone, shows a grid-mesh potential-temperature contour, a transparent moisture-mixing ratio surface, and trajectories over a topographical map. Figure 11 depicts underground water trajectories with transparent clay layers and text marking the locations of wells.

To enhance the depth information of our videotape animations, we usually combine time animation with a slight rocking of the three-dimensional scene.[3] The amplitude of this rocking is only one or two degrees and may be either vertical or horizontal, depending on the alignment of objects in the scene. This rocking is a very effective depth cue and is not ambiguous with most meteorological motions. We experimented with combining time animation with a steady rotation, but this created great ambiguity between the actual motion of objects and the apparent motion caused by rotation.

Transparency is an important part of our images, but it has presented some special problems. To render transparent surfaces efficiently, we developed a modified Z-buffer algorithm.[2] We analyzed and rendered each of the images in Figures 1-11 in about 30 seconds on our IBM 4381. Our workstation also poses a difficult problem for transparency: It does not separate red, green, and blue channels in its frame buffer. Rather, it was designed for false coloring of two-channel satellite data. We solved this problem by designing our images so that their pixels lie on two-dimensional planes through three-dimensional color space.

Interactive visualization

Our experience with very large data sets has shown the importance of interactivity for their visualization. Thus, we are currently developing a highly interactive version of our system—using the Stellar GS-1000 graphics supercomputer—to produce three-dimensional animations in real time (i.e., drawing at the animation rate). This system can provide an interactive window into data sets containing tens of millions of points produced by numerical models and remote sensing instruments.[4]

Figure 12 is an image produced by this system. It shows contour surfaces for sulfur dioxide (green), water vapor (red), potential temperature deviation (blue), and depth-cued wind trajectories. We can animate such images at five to 10 frames per second from a data set containing 10 different physical variables over 121 time steps in a 25 × 35 × 17-point three-dimensional grid. This data set contains 20 million grid points in a five-dimensional array.

System features. The user controls the display with a mouse, and the system gives immediate response to commands to

- rotate, zoom, and pan in three dimensions;
- select any combination of the scalar variables and wind trajectories; and
- start/stop time animation and single-step time forward or backward.

The user can also select new defining levels for contour surfaces using the mouse. The computation of new surfaces occurs at a rate of about two per second and is asynchronous with the animation rendering. The new surfaces replace the old in the animation as they are computed.

Our system holds an entire data set in main memory, along with polygon and vector lists to represent the surfaces and lines generated from gridded data. We use compressed data structures to maximize the size of the data sets we can visualize. The 20-million-point data set is visualized on a 64-megabyte system; a 128-megabyte system could be used for a 50-million-point data set.

The animation involves decompressing the polygon and vector lists, transforming them to a two-dimensional projection, and shading the pixels. Because each frame of the animation is generated as it is displayed, the user can instantly rotate and zoom by changing the transform function.

The user can also select the combination of viewed physical variables and control the time stepping by changing which polygon and vector lists are processed. The levels of contour surfaces are changed by computing new polygon sets from the gridded data, and this is done asynchronously with animation.

Scientific use. Interactive visualization is crucial to earth scientists, who must vary the way they look at a large data set according to its content. Controlling the combination of variables allows the user to examine specific cause-and-effect mechanisms. Coordinated hand-eye control of the view angle is a powerful way to understand three-dimensional geometries. Control of time stepping is needed to concentrate on particular events. Varying the defining levels of contour surfaces is important to quantitative understanding of data throughout a volume. We are working on an interactive mechanism allowing the user to move a plane through a three-dimensional volume, with contour lines of selected variables rendered on the plane. This mechanism will provide higher information density on the selected plane, so the user can concentrate on a specific spatial region.

A large data set often contains many interesting events and cause-and-effect links. Our system's interactivity lets scientists "play" with their data sets and pursue interesting physics at the pace of their own thoughts, without the distraction of waiting for the system to catch up.

Applications

Our visualizations are being used for teaching at the University of Wisconsin and other institutions. With Patricia Pauley of the University of Wisconsin Meteorology Department, we set up a series of animations on the binocular stereo workstation to help students in a weather laboratory better understand storm systems.[5] We are also making videotapes of weather visualizations available to other teaching institutions.

Scientists are using our visualizations to understand their data sets[6,7] and to present results at conferences. A videotape we produced with Louis Uccellini of the NASA Goddard Space Flight Center vividly depicts the development of the "Presidents' Day" storm that buried Washington, DC, in snow—when no snow was predicted.[8] This storm is being studied intensively by meteorologists trying to improve weather prediction. Our videotape clearly visualizes one simulation of this storm's dynamics and is being used for many conference and classroom presentations.

Scientists have begun using our interactive visualization system based on the Stellar GS-1000. Gregory Tripoli of the University of Wisconsin Meteorology Department used it for only a few hours, but that brief use helped him see and understand problems in his thunderstorm model that he had not seen during several years of viewing his data in two-dimensional plots. He has used this new understanding to correct problems in his model.

Clearly, interactive visualization is revolutionizing the ability of earth scientists to understand their enormous data sets. ☐

Acknowledgments

We would like to thank Verner Suomi, Frances Bretherton, Robert Fox, J.T. Young, and Robert Krauss of SSEC; James Dodge, Gregory Wilson, James Arnold and Paul Meyer of NASA; and Ian Reid, Craig Upson, David Kamins, Jeff Vroom, and Brian Kowalski of Stellar for their help and support. We would also like to thank the guest editor, Gregory M. Nielson, and the *Computer* referees for their suggestions.

This work was supported by NASA Marshall Space Flight Center (NAS8-33799 and NAS8-36292).

References

1. E. Smith, "The McIDAS System," *IEEE Trans. Geosci. Electron*, Vol. GE-13, July 1975, pp. 123-136.
2. W. Hibbard, "4D Display of Meteorological Data," *Proc. Workshop Interactive 3D Graphics*, ACM, New York, 1986, pp. 26-33.
3. W. Hibbard and D. Santek, "Presidents' Day Storm, Visualization/State of the Art: Update," *SIGGraph Video Rev.*, No. 35, 1988.
4. W. Hibbard, "A Next Generation McIDAS Workstation," *Proc. Conf. Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology*, Amer. Met. Soc., Boston, 1988, pp. 57-61.
5. P.M. Pauley, W. Hibbard, and D. Santek, "The Use of 4D Graphics in Teaching Synoptic Meteorology," *Proc. Conf. Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology*, Amer. Met. Soc., Boston, 1988, pp. 33-36.
6. D. Santek et al., "4D Techniques for Evaluation of Atmospheric Model Forecasts," *Proc. Digital Image Processing and Visual Communications Technologies in Meteorology*, SPIE, Bellingham, Wash., 1987, pp. 75-77.
7. P.J. Meyer and M.S. Seablom, "Application of the Four-Dimensional McIDAS to LAMPS Model Output," *Proc. Conf. Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology*, Amer. Met. Soc., Boston, 1988, pp. 33-36.
8. L. Uccellini, "Processes Contributing to the Rapid Development of Extratropical Cyclones," *Preprint Palmen Memorial Symp. on Extratropical Cyclones*, Amer. Met. Soc., Boston, 1988, pp. 110-115.

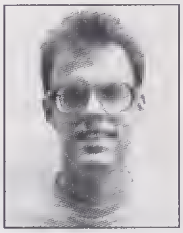

William L. Hibbard is a researcher at the Space Science and Engineering Center at the University of Wisconsin-Madison. His research interests are interactive computer graphics and image processing for earth sciences. He is a member of ACM, SIGGraph, SIGArt, and the IEEE Computer Society. Hibbard received a BA in mathematics and an MS in computer science from the University of Wisconsin-Madison in 1970 and 1974, respectively.

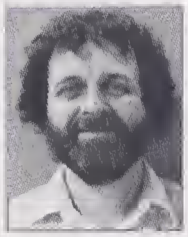

David A. Santek is a scientific applications programmer at the Space Science and Engineering Center at the University of Wisconsin-Madison. His interests include satellite data analysis, image analysis, and computer graphics. Santek received a BS in atmospheric and oceanic science from the University of Michigan in 1975 and an MS in meteorology from the University of Wisconsin in 1978.

Readers may contact the authors at the Space Science and Engineering Center, University of Wisconsin-Madison, 1225 W. Daytona St., Madison, Wis. 53706.

PRE-REGISTER FOR THE 14th Annual Conference on Local Computer Networks

Conference Site - The Radisson Plaza Hotel
35 South 7th Street
Minneapolis, MN 55402

Program Co-Chairs:

Larry Green
Protocol Engines, Inc.
1421 State Street
Santa Barbara, CA 93101
Phone (805) 965-0825
FAX (805) 564-8592

Marc Cohn
Northrop Corporation E593
8900 East Washington
Pico Rivera, CA 90660
Phone (213) 948-8979
FAX (213) 942-6478

General Chair:

Robert Linebarger
Computer Science Dept
242 TMCB
Brigham Young University
Provo, UT 84602
(801) 377-2616

Registration Information:

Pre-registration must be received on or before Thursday October 5, 1989 - no exceptions. We are not responsible for your delays, mail delays, etc. Registration fee for conference (Oct. 11-12) provides for conference proceedings, four refreshment breaks, two luncheons, and the banquet. Tutorial registration fee (Oct. 10) includes luncheon, tutorial notes, and two refreshment breaks. Make all checks payable to:
Conference on Local Computer Networks

Hotel:

A number of rooms has been set aside at the Radisson Plaza Hotel Minneapolis, 35 South 7th St., Minneapolis, MN 55402; (612) 339-4900. Please reference 14th Conference on Local Computer Networks.

Program Includes:

Tuesday, October 10, 1989

7:30 a.m. - 9:00 a.m. Registration
9:00 a.m. - 5:00 p.m. Tutorial 1
"Managing, Planning, & Controlling the Corporate Network," Jerry McDowell, McDowell Romero Inc.
9:00 a.m. - 5:00 p.m. Tutorial 2
"Practical Perspectives on OSI Applications," Chris Moore, The Wollongong Group

Wednesday, October 11, 1989

8:00 a.m. - 9:30 a.m. Registration
9:30 - 10:30 a.m. Keynote Session
"Building the 1990's Network," Howard Frank, Network Management Inc.
10:45 a.m. - 12:15 p.m. Technical Program
12:15 p.m. - 1:30 p.m. Lunch
1:30 p.m. - 5:00 p.m. Technical Program
5:30 p.m. - 6:30 p.m. Cocktail Party
6:30 - 9:00 p.m. Banquet
"The Magic of Hondo"

Thursday, October 12, 1989

8:30 a.m. - 12:15 p.m. Technical Program
12:15 p.m. - 1:30 p.m. Lunch
1:30 p.m. - 5:00 p.m. Technical Program

You are responsible for your own hotel reservations!

Registration:
Send registration and fee to:
14th Conference on
Local Computer Networks
Harvey A. Freeman
LANWORKS, Inc.
P.O. Box 506
Hopkins, MN 55343
(612) 591-5837

	Conference	Tutorial 1	Tutorial 2
Students (include copy of '89 paid fee statement)	$125 ____	$170 ____	$170 ____
IEEE Members	$235 ____	$170 ____	$170 ____
Non-Members	$295 ____	$215 ____	$215 ____

Payable in U.S. Dollars

Name ______________________________

Organization ______________________________

Address & Zip ______________________________

____

Phone ______________________________

Registrations received after October 5 will be returned. Registration after October 5 will be handled on-site at the rate of $280 for the conference and $205 for the tutorials for IEEE members and $355 for the conference and $260 for the tutorials for non-members. **Register early!**

B

14th Annual Conference On Local Computer Networks

October 10-12, 1989 Minneapolis, Minnesota

You are invited to attend

Tutorials - Panel Discussions - Paper Sessions - Discussion of Practical Experience

Technical Program

Wednesday, October 11, 1989

Track A Expert Panels	Track B Technical Papers	Track C Technical Papers
10:30-12:00 Management Chair: Bruce McClure, Synemetics Panel: Jerry MacDonald Howard Frank Jeff Case, Univ. of TN	10:30-12:00 FDDI Chair: Floyd Ross, UNISYS "FDDI Architecture," M. Teener, Apple "FDDI MAC Design Considerations," J. Torgerson, UNISYS "FDDI Connection Management," K. Ocheltree, IBM "The Role of Concentrators in FDDI Rings," J. Huchison, DEC	10:30-12:00 Network Performance I Chair: Alf Weaver, Univ. of VA "On Line Monitoring, The Role of Sampling in Network Measurements," P. Amer, L. Cassel, Univ. of DE "A Testbed for Token Ring Networks," H. Meyr, J. Tusch, Aachen Univ. of Technology "Using Multicast Communication for Dynamic Load Balancing in LANs," M. Ahamad, GA Inst. of Technology "Performance Analysis of a LAN-Based Remote Terminal Protocol," S. Uppal, HP
1:30-3:00 FDDI Chair: Floyd Ross, UNISYS Panel: Ken Ocheltree, IBM Bruce McClure, Synemetics Jerry Huchison, DEC	1:30-3:00 LAN Applications Chair: Bob Linebarger, BYU "LAN Cabling Design for the New BLS Building," P. Stevens, Bureau of Labor Statistics "Manageability, a Prime Consideration in Internetwork Design," M. Stevens, Boeing Computer Services "An RS-232-C Local Area Network," S. Leonard, T. Cochran, Atmel Corp. "A Behavior Model for the Application Services in Manufacturing," F. Lepage et al., Univ. of Nancy	1:30-3:00 Network Management Chair: Bruce McClure, Synemetics "FDDI Station Management," B. McClure, Synemetics "Network Management on Hughes Aircraft's Engineering Design LAN," P. Ho, Hughes "LAN Management Technology," S. Haugdahl, K. Thurber, Architecture Technology Corp. "Internet Management Using SNMP," J. Case, Univ. of TN
3:30-5:00 Fast Transport Protocols Chair: Alf Weaver, Univ. of VA Panel: Dick Watson, LLNL Larry Green, Protocol Engines, Inc. Carey Williamson, Stanford Kent Bossange, Ultra	3:30-5:00 MANs Chair: Jim Mollenauer, Prime "Fairness Issues of the DQDB Protocol," P. Martini, Lehrstubl for Informatik "Fairness Analysis & Technique to Extend Distance for 802.6," Y. Yaw, Synetics "IEEE 802.6 Update," J. Mollenauer, Prime "A Fair DQDY Access Control Mechanism," K. Kahalil, M. Koblentz, Bellcore	3:30-5:00 Emerging LAN Technology Chair: Ernst Biersack, Bellcore "Broadband ISDN Protocols and Interfaces," A. Hac, AT&T "Local Area Network Based on Fiber Optic Technology," K. DeMartino, Dynamics Research Corp. "Efficient Information Dissemination in Computer Networks," D. Agrawal, A. Malpani, UCSB "A New Buffer Insertion Ring," L. Kung et. al., National Cheng Kong University, ROC

Thursday, October 12, 1989

Track A Expert Panels	Track B Technical Papers	Track C Technical Papers
8:30-10:00 Standards Chair: Carl Cargill, DEC Panel: Floyd Ross, UNISYS Jim Mollenauer, Prime Buzz Rigsbee, Boeing	8:30-10:00 Distributed Networks Chair: Doug Jensen, Concurrent Computer "A Protocol for Program Migration," S. Ramakrishnan, Bowling Green Univ. "OSI Application Layer Standards Analysis for Distributed Applications," C. Nguyen, Universite Blaise Pascal "Task Allocation Schemes for Response Time Improvement," M. Kennedy, A. Badami, Case Western Reserve "Reliable Client-Server Communication in Distributed Programs," K. Ravindran, Indian Inst. of Science, S. Chanson, Univ. of British Columbia	8:30-10:00 Routing Chair: Jeff Yaplee, DEC "A Simple, Table-Free, Distributed Routing Algorithm for 802.6," Y. Yaw, Synetics "On the Graceful Degradation of Routing Given Insufficient Resources," P. Tsuchiya, Mitre "Integrated Routing," W. Seifert, Wellfleet "Building Multi-way Gateways for Local Area Networks," C. Yang, G. Xu, Univ. of North TX
10:30-12:00 Internetworking Chair: Bill Seifert, Wellfleet Panel: Jeff Yaplee, DEC John Hart, VITALINK Paul Tsuchiya, Mitre	10:30-12:00 Gigabit LANs Chair: Don Tolmie, Los Alamos National Labs "Multiple Crossbar Network," R. Hoebechheinrich, Loas Alsmos National Labs "How to Design a Gigabit LAN," H. Salwen, Proteon "Gigabit Networking in a Meshed HSC Environment," J. Hughes, Network Systems "Multiple Crossbar Network Interface," W. Hedberg, DEC	10:30-12:00 Networking Performance II Chair: Alf Weaver, Univ. of VA "Performance Measurements of a Communications Protocol Implemented for Distributed, Real-Time Systems," A. Van Der Wal, G.J.W. Dijk, Eindhoven Univ. of Technology "Explicit Versus On-Demand Connection Setup," C. Williamson, Stanford, Univ. "Measured Ethernet Performance for Multiple, Large, File Transfer," S. Francis et al., Univ. of KA "Performance Analysis of FDDI," R. Sanker, Y. Yank, Univ. of South FL
1:30-3:00 Gigabit LANs Chair: Bob Grow, Cumulus Panel: Don Tolmie, Los Alamos Howard Salwen, Proteon Jim Hughes, NSC	1:30-3:00 FDDI Networks Chair: Floyd Ross, UNISYS "NASA Space Station LAN," A. Weaver, T. Strayer, Univ. of VA "SAFENET: A LAN for Navy Mission Critical Systems," D. Green, D. Marlow, NSWC "The FDDI Link Error Rate Monitor," G. Kajos, D. Hunt, Prime "The Fiber Distributed Data Interface," J. Figueroa, AT&T	1:30-3:00 Network Reliability Chair: Tim Strayer, University of VA "An Argument for Judicious Use of Error Detection Mechanisms," T. Strayer, Univ. of VA "Performance Modeling and Analysis of a Degradable Computer Communications Network with Blocking," J. Hsieh, D. Ucci, IL Inst. of Technology "An Availability Model for Distributed Computing Environments," H. Mutlu, AT&T "Network Management with OSI and POSIX," M. Colburn, Open Systems Architecture
3:30-5:00 15th LCN Planning Chair: Harvey Freeman, LANWORKS	3:30-5:00 Fast Protocols Chair: Dick Watson, Lawrence Livermore Labs "The Delta-T Transport Protocol," D. Watson, Lawrence Livermore National Labs. "The Xpress Transfer Protocol," A. Whaley, Protocol Engines Inc. "An Overview of the VMTP Transport Protocol," C. Williamson, Stanford Univ. "High Speed TP-4," K. Bossange, ULTRA	3:30-5:00 Network Availability Chair: Oliver IBE, DEC "Functional Availability of Computer Communications Networks," A. Shah, R. Roy, CSC "Performance of Playthrough Rings with Mixed Types of Traffic," C. Silio, Univ. of MD "An Orientation for Production Quality Networks," J. Rao, McDonnell Douglas "Local Network Interconnection Schemes & Their Availability Analysis," O. Ibe, R. Howe, DEC

Sponsored by:

IEEE COMPUTER SOCIETY

TC-Computer Communications

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

IEEE

The Role of Visualization in the Simulation of Quantum Electronic Transport in Semiconductors

Norman C. Kluksdahl, Alfred M. Kriman, and David K. Ferry
Arizona State University

Using the Wigner function to simulate semiconductor processes, then depicting the results graphically, can lead to new insights concerning these devices.

In simulating a semiconductor device, examining the raw numerical data rarely fosters an understanding of physical processes within the device. Visualization tools provide a useful method for interpreting this voluminous data and allow a better understanding of the semiconductor. Pictorial representation of data can reveal results not readily discerned from the data alone. Furthermore, animating a sequence of data frames permits visualization of time-evolving systems.

These advantages hold especially true when studying ultrasmall devices such as the resonant tunneling diode. The RTD is a gallium-arsenide semiconductor structure in which two thin layers (a few nanometers thick) of aluminum gallium arsenide (AlGaAs) form quantum tunneling barriers around a thin layer of gallium arsenide. The quantum tunneling probability in this structure possesses a sharp peak (Figure 1), corresponding to resonant tunneling.[1] Experimental devices incorporating this structure show negative differential conductivity in their current-voltage (I-V) characteristics,[2] as predicted by theory.[1,3]

Efforts to model the RTD met with only limited success[3] until employment of the Wigner function formalism.[4-6] Coupled with increased graphics capability, this has led to a better understanding of the RTD and to the prediction of new physical phenomena within the RTD such as plasma oscillations during transient switching.

The Wigner function

A classical system often employs a phase-space distribution function. This approach, the foundation of Monte Carlo simulations,[7] has found frequent use in modeling large semiconductor devices. It fails, however, when the feature sizes are such that the internal potential varies over a length shorter than a characteristic electron wavelength.[4]

0018-9162/89/0800-0060$01.00 ©1989 IEEE

An alternative approach, the Wigner function,[8] uses a formalism that correctly incorporates all the quantum mechanics inherent in a problem. The Wigner function is the Fourier transform of the density matrix, taken with respect to a difference variable:

$$f_w(x,p) = \frac{1}{2\pi\hbar}\int dy\, e^{ipy/\hbar} \varrho(x+\frac{y}{2}, x-\frac{y}{2}) \quad (1)$$

where $\varrho(x,x')$ is the density matrix describing how the system is composed of a mixture of the system's eigenstates. This transformation, the Wigner-Weyl transform, when applied to the density matrix time-evolution equation, gives the equation for the Wigner function:

$$0 = \frac{\delta f}{\delta t} + v\frac{\delta f}{\delta x} + \frac{1}{\hbar}\int dP \int dy\, f(x,p+P) \sin(2Py) * [V(x+y) - V(x-y)] \quad (2)$$

The Wigner function has many advantages.[4,8] It is a real function, unlike the complex density matrix from which it is derived. The equation of motion is similar to the classical Boltzmann equation. And, finally, expectation values of operators such as the current operator are easily evaluated. This formalism does, however, have a few quirks.[4,8] The Wigner function is not a true probability distribution function in that it is not positively semidefinite—a direct consequence of the Heisenberg uncertainty principle. In areas where quantum mechanical effects dominate, the distribution may manifest negativity. This necessitates caution in evaluating expectation values to ensure meaningful results. More importantly, quantum mechanically exact boundary conditions must be specified to avoid excluding some orders of the quantum corrections.

The time-evolving Wigner function is found by numerically solving the equation of motion (Equation 2) with explicit finite differencing. Good convergence is obtained by retaining the second-order term in the Taylor series expansion. The final form of the equation of motion, after some algebra, is

$$f(t+\Delta t) = f(t) - \Delta t[v\frac{\delta f}{\delta x} + \theta(f)] + \frac{\Delta t^2}{2}[v^2\frac{\delta^2 f}{\delta x^2} + v\frac{\delta\theta(f)}{\delta x} - \frac{\delta\theta(f)}{\delta t}] \quad (3)$$

(In this notation, θ is the operator for the integral forcing function.) Also, measures must be taken to ensure that the nonlocal momentum operator in θ does not introduce artificial numerical reflections.[4,9]

The resulting discretized mesh can be thought of as a collection of equations, each of which occupies a "slice" of the mesh with a fixed momentum value p. These slices are coupled through the θ operator of Equation 3. Each slice then has a characteristic velocity ($v=p/m$) and direction of information propagation, and a corresponding boundary at which information enters. This is the specified boundary, then, for the slice.

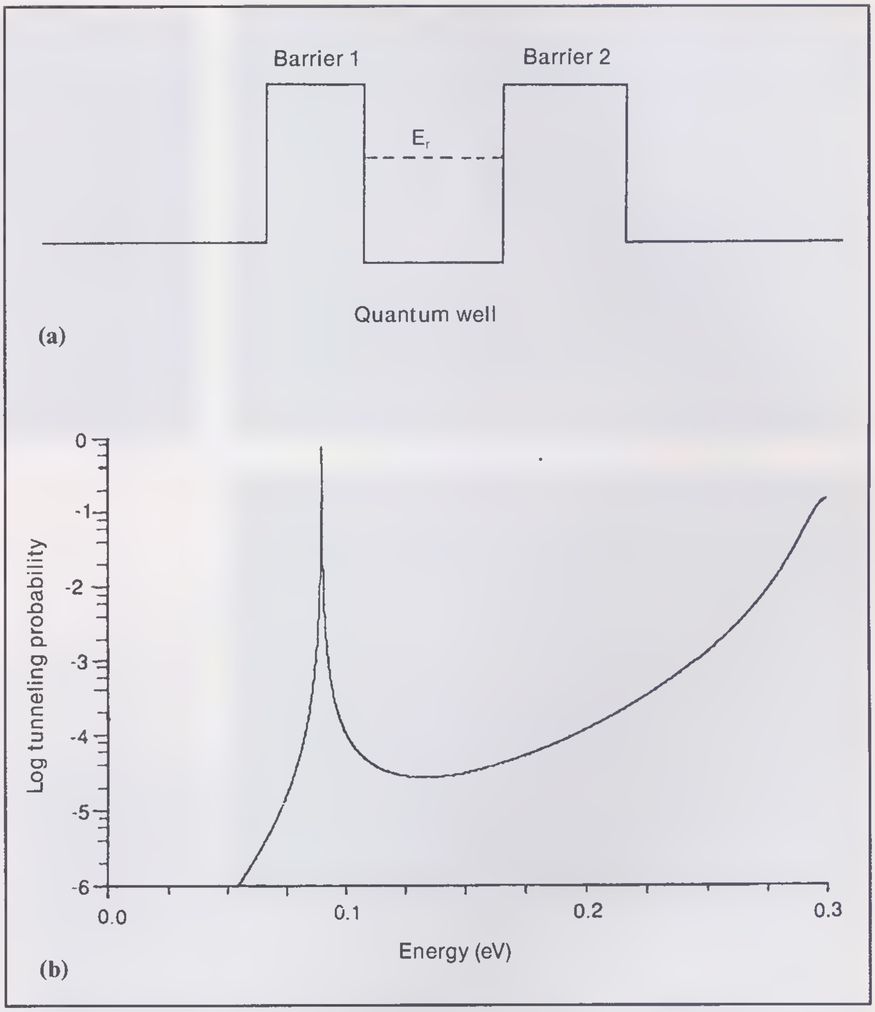

Figure 1. A resonant tunneling structure: (a) the conduction band; (b) the calculated tunneling probability for a typical resonant tunneling structure. In a device, the barriers are $Al_xGa_{1-x}As$, while the well and the contact regions consist of GaAs. E_r is the resonant energy level.

Tunneling problem

The simplest systems that can be studied are wave packets tunneling through a quantum barrier. Such examples, rife in elementary quantum mechanics textbooks,[1] provide a method of verifying the accuracy of the Wigner function equation of motion. Because the Wigner function shows both position and momentum, it offers new insights into the processes of tunneling and reflection, and not mere repetitions of already-solved problems. When graphically represented, the Wigner function displays a wealth of information about tunneling wave packets.

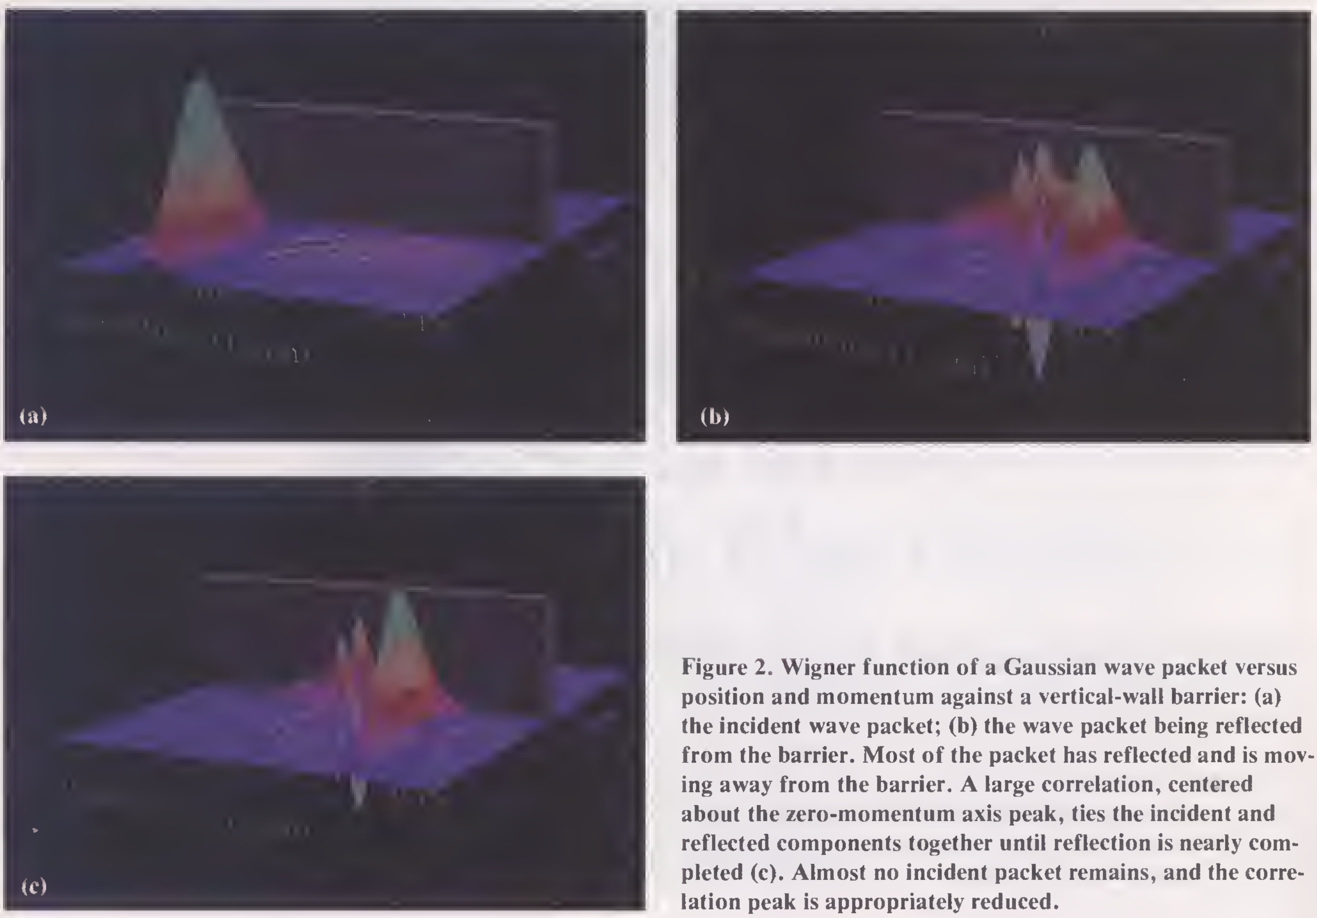

Figure 2. Wigner function of a Gaussian wave packet versus position and momentum against a vertical-wall barrier: (a) the incident wave packet; (b) the wave packet being reflected from the barrier. Most of the packet has reflected and is moving away from the barrier. A large correlation, centered about the zero-momentum axis peak, ties the incident and reflected components together until reflection is nearly completed (c). Almost no incident packet remains, and the correlation peak is appropriately reduced.

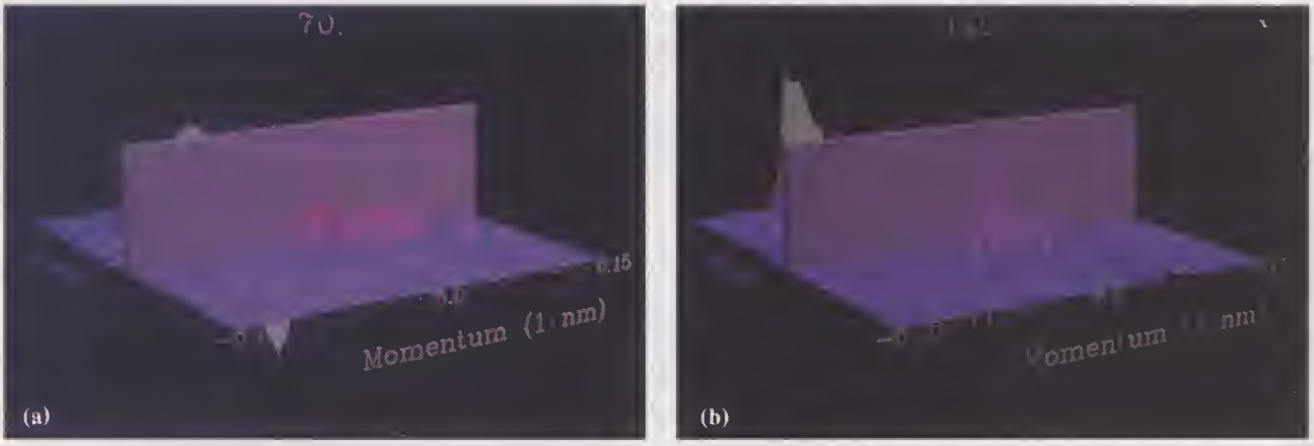

Figure 3. The tunneling wave packet, viewed from the opposite side of the barrier: (a) the transmitted wave packet is against the barrier, with two components—the actual wave packet on the right, and a correlation peak on the left; (b) the tunneling is completed, and most of the transmitted packet has left the domain. The correlation is a series of rapid oscillations at the zero-momentum axis.

A Gaussian wave packet incident on a finite-thickness quantum barrier will split into two components: One fraction will tunnel through the barrier, and the remainder of the wave packet will be reflected. The Wigner representation of the wave packet has a Gaussian form in both position and momentum.

A quantum barrier potential is inserted into θ, an initial wave packet is placed on the mesh, and Equation 3 is iterated forward in time. As the iteration proceeds, the wave packet is plotted versus position and momentum, with amplitude represented by both height and color. Figure 2a shows the wave packet shortly after $t = 0$. Points with positive momentum move into the page; negative momentum corresponds to motion out of the page. In Figure 2b, the wave packet is in the late stages of reflecting; a Gaussian wave packet with negative momentum is forming. Some of the incident packet is visible at the barrier, and correlation between the incident and reflected wave packets is manifested as peaks in the distribution around the zero-momentum axis. In Figure 2c, the reflection is nearly completed: The incident wave packet has nearly vanished, the reflected wave packet has grown, and the correlation has been reduced. The negative peaks are a consequence of the uncertainty principle.

These plots, and the animated representation of the tunneling process, reveal that the effect of a barrier extends outward from its physical location and turns the wave packet slowly. Quantum reflection takes a finite time. In contrast, the classical view of this problem yields an impulse force and instantaneous reflection (for nondeformable bodies).

The tunneling problem can be viewed from the other side of the barrier, as in Figures 3a and 3b. In Figure 3a, the transmitted wave packet shows as two parts: The rightmost wave packet at the barrier, and a correlation peak centered about the zero-momentum axis. Figure 3b shows the transmitted wave packet after the tunneling has completed. The transmitted wave packet is barely visible at the edge of the simulation domain. The most prominent feature of this figure is that the correlation appears as rapid oscillations along the zero-momentum axis. These oscillations do not contribute to density, but contain information about the correlation between the reflected and transmitted wave packets. When this correlation is retained, the system can be reversed to recover the initial wave packet. When it is removed, the system behaves as if two separate, uncorrelated wave packets were striking the tunnel barrier.

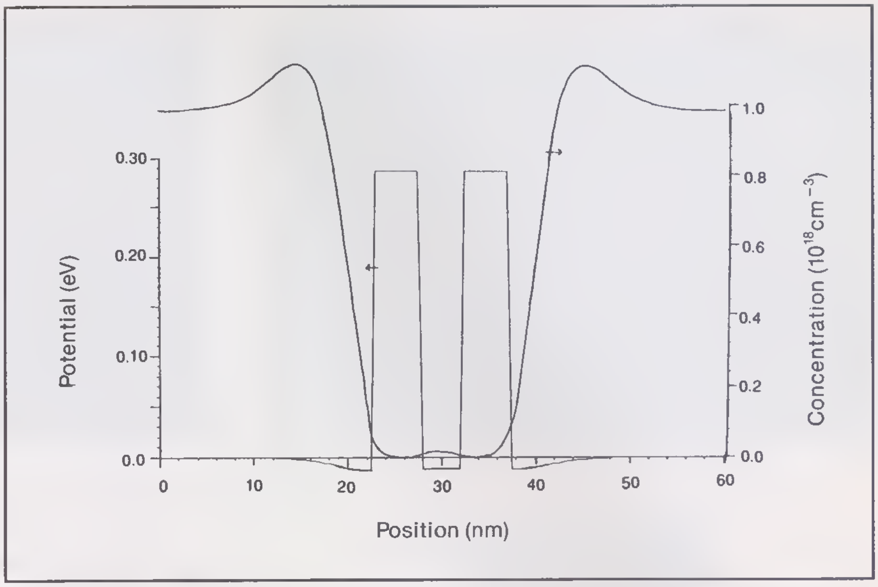

Figure 4. The equilibrium self-consistent conduction band and the electron density of the modeled resonant tunneling diode (RTD).

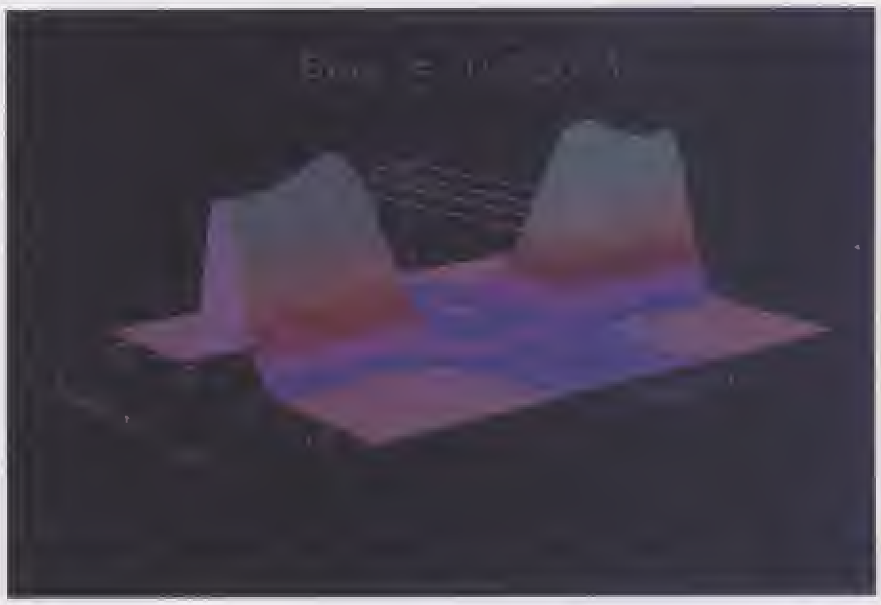

Figure 5. The equilibrium Wigner distribution for the RTD, versus position and momentum. The quantum barriers are indicated by the white wire frames.

Device simulations

The problem of the resonant tunneling diode is solved by first finding a density matrix and potential that are self-consistent.[4] The potential depends upon the densities of electrons and ionized donors in the device; the ionized donor concentration and the electron density depend on the potential. The structure of the device is shown in Figure 1a, with the barriers and well all 5 nanometers thick. Figure 4 shows the equilibrium self-consistent conduction band and the electron density.

The equilibrium density matrix, transformed into the Wigner function representation, becomes the initial condition for the time-evolution equation (Equation 3). Figure 5 shows the equilibrium Wigner function as density versus position and momentum. The white wire frames indicate the positions of the quantum barriers. This distribution shows two interesting features: quantum repulsion from the bar-

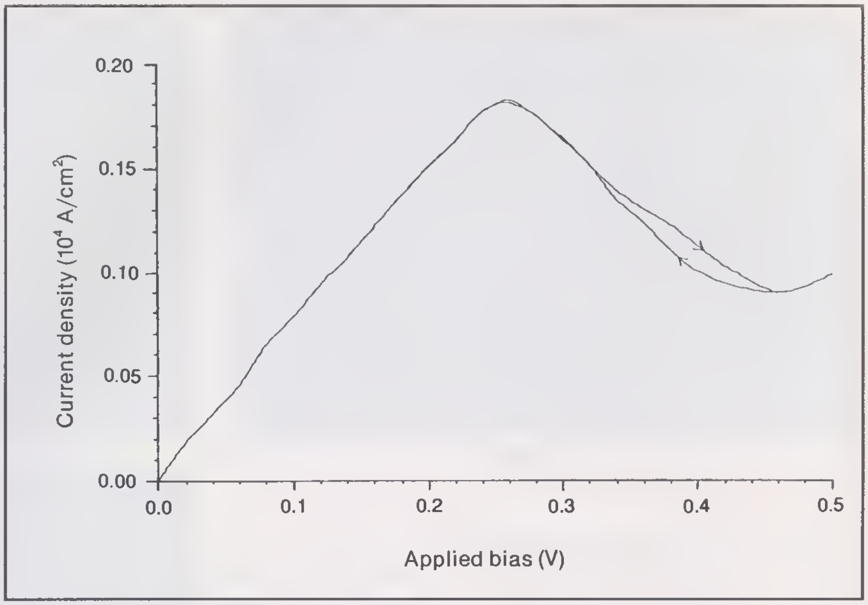

Figure 6. The calculated current-voltage curve of the modeled RTD.

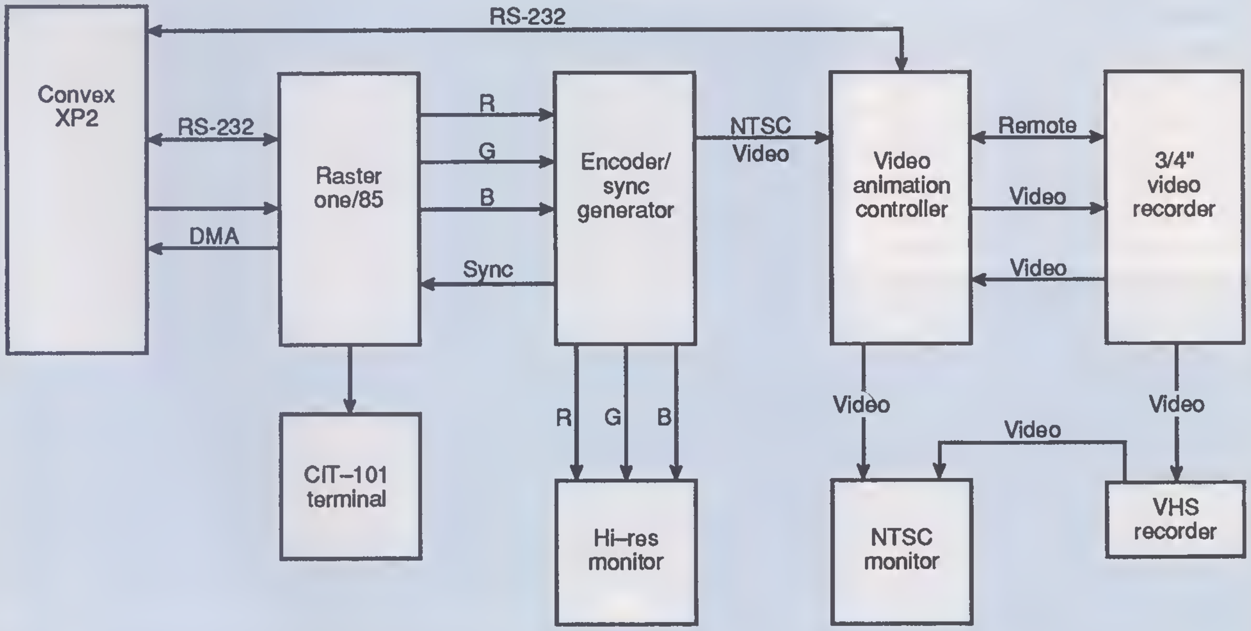

Block diagram of the animation equipment. Animation is totally automated under control of the user program on the Convex mainframe.

Animation studio

All computations are carried out on a Convex mainframe. Graphics are also generated on the Convex and then piped to a graphics workstation. An encoder/sync generator generates National Television System Committee video, providing source data for an animation controller. The controller is driven by the Convex and synchronizes video frames onto a ¾" video tape recorder (see figure).

The application program generates the images through a graphics library, then hands control to an animation library. After the image is recorded, the animation library passes control back to the calling program, and the system is readied for another image. Using this approach, videotape generation is a hands-off automated process. Videotapes are finished by using an editing controller, a second ¾" video tape recorder, and a title generator. The ¾" master tape can then be copied onto VHS tape for distribution or presentation.

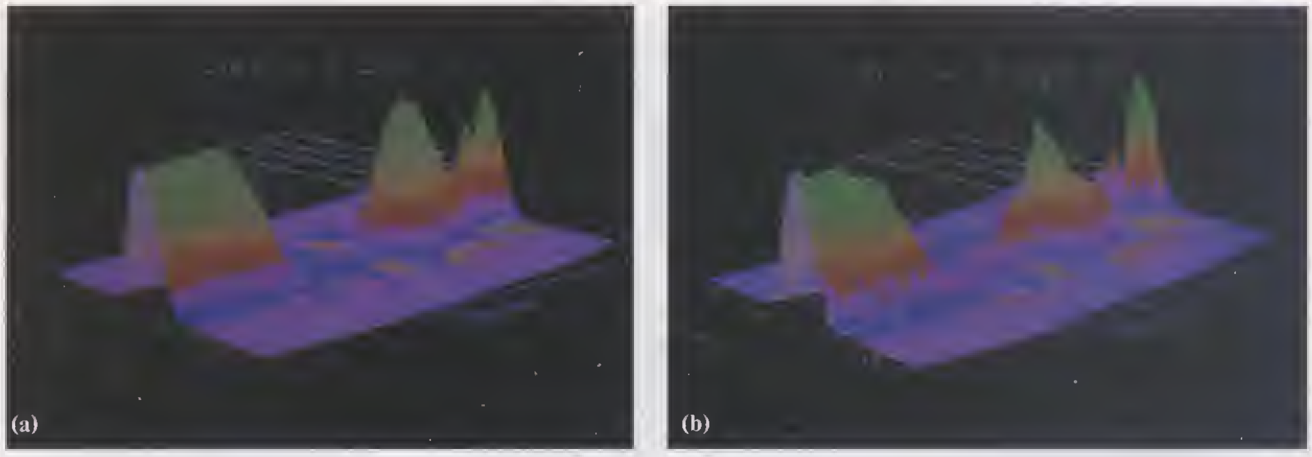

Figure 7. Two different RTD distribution profiles, varied through applied bias: (a) the RTD Wigner distribution for an applied bias of 0.22 volts shows depletion at the cathode contact, the anode accumulation swept out, and tunneling components of current; (b) the RTD distribution for an applied bias of 0.44 volts with the cathode end of the device dominated by a triangular potential well, resulting in a quantized state that is coupled to the anode end by tunneling through the barriers.

riers leading to high momentum tails in the distribution near the barriers, and an accumulation near the barriers caused by the self-consistent potential.

A small potential is added to one end of the device, and the system is time-evolved by alternately stepping (Equation 3) in time and solving Poisson's equation until steady-state conditions are achieved. The current is calculated, and the process is repeated until the total I-V curve has been calculated (Figure 6).

Figure 7a shows the steady-state distribution for an applied potential of 0.22 volt, corresponding to the peak of the I-V curve. The base of the distribution is raised proportionally to the self-consistent potential $V(x)$ (but not on the same scale as the wire-frame barriers). The accumulation layer in the left contact region, the anode, has been swept out. The distribution in the right contact region, the cathode, is heavily depleted between the boundary and the barrier. At the cathode contact, the number of electrons leaving the device has been strongly reduced. Figure 7b shows the distribution for an applied bias of 0.44 volt, the valley of the I-V curve. Very few electrons are leaving the cathode contact. The cathode region has a triangular potential well between the boundary and the potential barrier, and the distribution at this point is a ring, the expected form of a quasibound state. Most notable is the strong tunneling through the barriers between the contact regions. This tunneling is symmetric with respect to momentum and represents coupling between the anode and the quasibound state in the cathode well.

Transients have been studied by abruptly switching the bias from steady-state conditions at the peak of the I-V curve to the valley. A plot of each time step is synchronously edited onto videotape so that the resulting animation shows the transient distributions. As the device switches, the charge between the barriers is forced out, altering the potential. As the potential changes, the tunneling probability is changed, and the current changes correspondingly. A mechanism is thus established for oscillatory transient currents, and these have been observed.[5] One newly modeled result reveals these current oscillations to be sufficiently strong to cause large-scale collective oscillations within the entire device. Only with animation is this behavior observed in the model.

Graphical representation of wave packets during quantum tunneling shows explicitly the correlations between the transmitted and reflected components. When the time-evolving wave packet is animated, details of the reflection and transmission processes become strikingly evident.

In device simulation, a phase-space distribution is particularly advantageous, but such distributions require powerful visualization techniques in order to obtain a better understanding of the behavior predicted by models and to predict entirely new phenomena. Animated sequences show time-dependent phenomena which cannot be seen in single-image representations. ☐

References

1. D. Bohm, *Quantum Theory*, Prentice Hall, Englewood Cliffs, N.J., 1951.
2. T.C.L.G. Sollner et al., "Resonant Tunneling through Quantum Wells at Frequencies up to 2.5 THz," *Appl. Phys. Lett.*, Vol. 43, Sept. 1983, pp. 588-590.
3. B. Ricco and M.Ya. Azbel, "Physics of Resonant Tunneling: The One-Dimensional Double Barrier Case," *Phys. Rev. B*, Vol. 29, Feb. 1984, pp. 1,970-1,981.
4. N.C. Kluksdahl, A.M. Kriman, and D.K. Ferry, "Wigner Function Simulation of Quantum Tunneling," *Superlattices and Microstructures*, Vol. 4, No. 2, 1988, pp. 127-131.
5. N.C. Kluksdahl et al., "Transient Switching Behavior of the Resonant Tunneling Diode," *IEEE Elec. Dev. Lett.*, Vol. 9, Sept. 1988, pp. 457-459.
6. W.R. Frensley, "Wigner Function Model of a Resonant Tunneling Semiconductor Device," *Phys. Rev. B*, Vol. 36, July 1987, pp. 1,570-1,580.
7. C. Jacoboni and L. Reggiani, "The Monte Carlo Method for the Solution of Charge Transport in Semiconductors with Application to Covalent Materials," *Rev. Mod. Phys.*, Vol. 55, July 1983, pp. 645-705.
8. P. Carruthers and F. Zachariason, "Quantum Collision Theory with Phase-Space Distributions," *Rev. Mod. Phys.*, Vol. 55, Jan. 1983, pp. 245-285.
9. B. Engquist and A. Majda, "Absorbing Boundary Conditions for the Numerical Simulation of Waves," *Math. Comp.*, Vol. 31, July 1977, pp. 629-651.

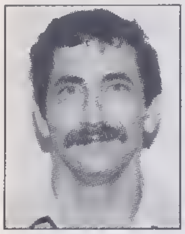

Norman C. Kluksdahl is a faculty associate in the Center for Solid-State Electronics Research at Arizona State University. He is currently studying quantum transport in semiconductors, although his interests also extend to VLSI and computer architecture. Kluksdahl served as technical consultant on the purchase of minisupercomputer and graphics equipment for the center and worked to establish a computer animation laboratory in the center.

Kluksdahl is a member of IEEE and Sigma Xi. He was a 1977 Presidential Scholar and a recipient of Motorola and Honeywell scholarships. Kluksdahl received his BS and MS in computer science from South Dakota School of Mines and Technology in 1981 and 1984, respectively, and his PhD in electrical engineering from Arizona State University in 1988.

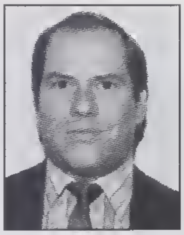

Alfred M. Kriman is a senior research analyst in the Center for Solid-State Electronics Research at Arizona State University. From 1984 to 1986 he was a National Research Council postdoctoral fellow at the Naval Research Laboratory in Washington, DC. Kriman's research interests currently center on the analysis of the electronic structure and properties of microstructures.

Kriman is a member of the IEEE, the American Physical Society, Phi Beta Kappa, and Sigma Xi. He received an AB from Rutgers College in 1978, and AM and PhD degrees from Princeton University in 1980 and 1984, respectively, all in physics.

David K. Ferry assumed his current position as chair of the Electrical Engineering Department of Arizona State University this year. He had been director of the Center for Solid-State Electronics Research and professor of electrical and computer engineering at the university since 1983. Previously, he was with Colorado State University and the Office of Naval Research. He served on the NAS/NRC study on Thin-Film Microstructure Science and Technology, has codirected two NATO Advanced Study Institutes, and has served on program and organizing committees for several conferences. Since 1982, he has been a member of DARPA's Materials Research Council. He is author or coauthor of some 200 scientific works.

Ferry is a fellow of the American Physical Society. He received BS and MS degrees from Texas Tech University, Lubbock, in 1962 and 1963, respectively, and a PhD degree from the University of Texas, Austin, in 1966.

The authors' address is Center for Solid-State Electronics Research, Arizona State University, Tempe, AZ 85287-6206.

When it comes to choosing compilers, MetaWare is the right choice!

Developing the next generation of software products is serious business. You need the best tools to produce the best code. You need **MetaWare**.

Superior Compilers

A compiler that handles large programs with ease, while producing the expected results, is the key to developing the highest quality applications. Many of **MetaWare's** customers say that **High C** ™ and **Professional Pascal** ™ are the highest quality compilers in the industry. They are reliable, and well documented, and their superior diagnostic messages help to produce better products more quickly. **There are no surprises with MetaWare compilers.**

Compiler Features

ANSI Standard with extensions • Generates small, fast executables • Support of 80x87, Weitek 1167/3167, 68881, and Am29027 math co-processors • Global common subexpression elimination • Live / dead code analysis • Constant propagation, copy propagation • Tail merging (cross jumping) • And many more!

Multiple Platform Support

MetaWare uses its own **Translator Writing System** to create all of its compiler products. Common components and library functions are shared across the product line. Improvements to compilers are quickly realized on all platforms.

MetaWare, High C and Professional Pascal are trademarks of MetaWare Incorporated; UNIX is a trademark of AT&T; Other products mentioned are trademarks of the respective companies indicated.

Supported Platforms (re: Dhrystones)

• Sun-3	— >50% > resident compiler.
• Sun 386*i*	— >50% > resident compiler.
• Sun-4	— >25% > resident compiler.
• PC: DOS, OS/2	— 3-10% > Microsoft C; 30%>MS Pascal, Lattice C.
• 386 32-bit DOS	— No real competition.
• 286, 386 UNIX	— 66% > than pcc on 386.
• VAX VMS	— ≈ DEC's excellent C and Pascal; Host for cross compilers and TWS, not Native.
• VAX Ultrix	— 19% > pcc on Dhrystone; much > Berkeley Pascal.
• RT PC	— 93% > 4.3bsd port of pcc.
• AIX/370	— Much better than any 370 C and VS Pascal.
• AMD 29K	— >40,000 Dhrystones.
• Intel i860	— >70,000 Dhrystones at 33 MHz.

So when it comes to selecting a compiler company, **you need the best!**

Can you really afford anything less?

MetaWare™
INCORPORATED

The Compiler Products for
Professional Software Developers

2161 Delaware Avenue • Santa Cruz, CA 95060-5706
Phone: (408) 429-6382 • FAX: (408) 429-9273

Reader Service Number 8

The key to new doors in Science & Technology.

The NYNEX Science & Technology Center is the catalyst in our commitment to information industry leadership. The Center is home to a diverse family of laboratories working in specific, yet related areas of research—all striving to develop solutions that will meet the varied and changing needs of today's information age customer. If you have an advanced degree in Electrical Engineering, Computer Science, Physics, Mathematics, or Psychology, we have the key to exceptional challenge in one of the following areas:

Applied Research focuses on the development of new service and network concepts to meet current and future customer needs for the movement, retrieval and presentation of information.

Advanced Technology Development evaluates and develops prototype state-of-the-art services and hardware/software products, as well as provides technical support to the entire NYNEX family of companies.

Whatever area you choose, you can shape the development of real-world communications and information delivery through both marketing and technology.

NYNEX is firmly committed to a policy of providing equal employment opportunity and affirmative action. For further information and consideration, contact:

NYNEX Science & Technology Center
500 Westchester Ave., Room 1A4
White Plains, NY 10604

Looking for options? Looking for opportunities? The answer is

NYNEX

Scientific Visualization at Research Laboratories

Large scientific research centers are developing visualization methodologies and tools. This special section examines work under way at Los Alamos, Lawrence Livermore, the National Center for Supercomputing Applications, NASA Ames, and the Naval Research Laboratory.

Lawrence J. Rosenblum, Naval Research Laboratory

The concept of *scientific visualization* reaches back into prehistoric times when a caveman drew a map of his local environment on his cave wall. In antiquity, legend tells us that Archimedes was slain by a Roman soldier while visualizing figures sketched in the sand. In this century, chemists began to understand the structure of matter and satisfied the need to visualize molecules with wooden and plastic models. Visualizing data and concepts is not new, nor is it computer dependent.

In the computer age, we have progressed through line-printer output, contour plots, etc., to more sophisticated techniques. Yet, scientific visualization has only emerged as a technology in the last two or three years.

Visualization is no longer a useful supplement to scientific and engineering progress; it is now essential for many disciplines. Complex computational models, executing on supercomputers, produce simulated data in quantities that require geometrically based algorithms for interpretation. From multibeam echosounders mapping the ocean bottom to radio telescopes studying the heavens, sensors provide more numerical data than can ever be examined point by point. Large scientific research centers have responded by establishing and/or expanding groups devoted to the development of visualization methodologies and tools for an integrated understanding of very large data sets.

There are three levels of visualizing data for numerical simulations: *postprocessing, tracking,* and *steering*. Postprocessing refers to portraying the data after the numerical simulation is complete. Tracking provides real-time displays, allowing termination of a faulty simulation. Steering permits interaction with the computation; simulation parameter values can be altered midstream and the computation continued.

To date, most efforts deal with techniques for effective postprocessing: methods for transferring simulation data to the user and then displaying and interactively manipulating it. Steering offers great potential benefits, but both hardware and toolset problems must be solved. Steering is not new, but the complexity of many of today's simulations enhances its value. While emphasis is being placed on supercomputer steering, perhaps the "graphics minisupercomputers," with their tightly integrated computational and graphical power, are the platform for learning how to effectively steer complex simulations.

Current visualization activities

The articles in this section of the special issue examine aspects of scientific visualization work at five major laboratories and centers in the United States. With a half dozen or so supercomputers each—the numbers change as older models are replaced—Lawrence Livermore National Laboratory (LLNL) and Los Alamos National Laboratory (LANL) each have among the largest concentrations of supercomputing power in the world. "Distributed Visualization at Los Alamos National Laboratory" by Richard L. Phillips describes several approaches to visualizing simulated data. The video animation facility of Karl-Heinz Winkler and associates at LANL represents the state of the art in scientific visualization, but is too expensive to replicate widely. Philips also presents methods used at LANL for connecting supercomputers with display units, ranging from dumb terminals to powerful workstations. Of particular interest is the Scientific Visualization Workbench, an inexpensive and effective way of using "televisualization" to move data from supercomputer to user and to modify the user's display.

"Visualization Tools at Lawrence Livermore National Laboratory" by Brian Cabral and Carol L. Hunter describes three projects at LLNL. The Magic Project is developing techniques to move data from a Cray and effectively display it on high-resolution color terminals in the scientist's office. Laser disk technology provides a real-time animation capability. The Graphics Workstation Project has examined visualization applications of workstations with graphics accelerators and is currently extending this work to "graphics minisupercomputers." The Advanced Visualization Research Project uses high-end graphics devices to perform algorithm development, especially in the

0018-9162/89/0800-0068$01.00 ©1989 IEEE

For the Serious Computer Graphics Professional:

IEEE
Computer Graphics
and applications

Subscribe Today!

☑ **Yes,** sign me up to receive *IEEE CG&A!*

☐ **I'm a member of an IEEE society, so I am paying the member price of $19 for a full-year subscription or if subscribing between March 1 and August 31, I can receive a half-year for only $9.50.**

☐ **I'm not a member of an IEEE society, but I am a member of another professional society, so I am paying the sister society price of $35 for a full-year subscription.**

NAME ____________________
ADDRESS ____________________
CITY ____________ STATE ____________ ZIP ____________
SOCIETY/ORGANIZATION ____________________
MEMBER NUMBER ____________________

☐ Check or money order payable to IEEE enclosed.

Charge my: ☐ MasterCard ☐ VISA ☐ American Express

Card number: ____________________ ________ Expires

Exp. Date 12/31/89

Signature: ____________________

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 38 LOS ALAMITOS, CA

POSTAGE WILL BE PAID BY ADDRESSEE

IEEE Computer Society
10662 Los Vaqueros Circle
Los Alamitos, CA 90720–9805

implementation of volumetric methods.

The National Center for Supercomputing Applications (NCSA), one of five National Science Foundation-sponsored supercomputer centers, serves a national user base rather than on-site users. In "Scientific Visualization and the Rivers Project at the National Center for Supercomputing Applications," Robert B. Haber describes NCSA's approach to the development of visualization tools for scientists who may be hundreds or thousands of miles away, as well as meeting the needs of researchers with special, complex requirements. Haber's Rivers (Research on Interactive Visual Environments) Project is an interesting and potentially important attempt to extend high-end visualization to the interactive steering of supercomputer simulations. Other groups at NCSA deal with "routine" graphics, with custom high-quality animations for scientists, and with the development of toolsets ranging from generating color palettes to volume visualization (see below).

At the Ames Research Center of the National Aeronautics and Space Administration, scientific visualization is more focused. "Scientific Visualization in Computational Aerodynamics at NASA Ames Research Center" by Gordon Bancroft et al. describes tools specifically oriented toward flow analysis problems. This approach is oriented toward performing graphics calculations on the supercomputer and using high-end workstations as rendering engines to display flow fields generated by aircraft flight. A suite of software provides both display and interactive analysis capability. The authors also describe animation hardware featuring Winchester disk technology, which allows the storage of more than one gigabyte of data and supports editing, special effects, and video output.

In addition to in-house computational science, the Naval Research Laboratory (NRL) supports a varied experimental research program. NRL scientists travel worldwide to perform experiments using platforms that include satellites, aircraft, ships, and submarines. "Visualization of Experimental Data at the Naval Research Laboratory" by Lawrence J. Rosenblum deals with the role of scientific visualization in understanding experimental data sets. The difficulty in extracting parameters of interest from the measurements is described for five diverse application areas, and the role of visualization in developing understanding of these data is illustrated. The role digital imaging and signal processing play should be noted. These areas are perhaps more difficult to generalize into toolsets than computer graphics, but they are an important part of the suite of visualization technologies.

Trends and directions

These articles make it clear that the postprocessing problem is well on the way to solution. Higher bandwidth networks are needed, as are more capable workstations, higher resolution, better interaction, etc., but these will arrive. The scientist will sit at his desk and see simulation results quickly and effectively. Effective steering tools are further away, and efforts such as the Rivers Project are important because successful steering tools will change the very nature of computational science.

One much-discussed issue in the scientific visualization community is "How much graphics processing should be done on the supercomputer and how much should be moved to the workstation?" A variety of approaches is seen. At NASA/Ames, the workstation is used primarily for rendering the aircraft in flight. At NCSA, data is downloaded from the supercomputer onto a minisupercomputer devoted to visualization processing. Meanwhile, LANL and LLNL are experimenting with a variety of techniques for moving data in unprocessed, partially processed, and graphical formats between supercomputer and display units that range from dumb terminals to high-end graphics workstations. Perhaps this issue is unresolvable because, in addition to the technical arguments, many factors in such a decision relate to organizational structure. The questions include: "Is time available on the supercomputer and at what cost?" and "How is the visualization developer's time paid?" The distinction should be drawn, however, between the need to develop general visualization tools, which require minimal interaction with the scientist, and the need for interdisciplinary teams of scientists and visualization specialists working closely together to attack front-line scientific issues. Examples of both can be found in these articles, and both are essential.

An illustration of this distinction can be seen in an important methodology called volume visualization (or volume rendering), mentioned in several of these articles. Because volume visualization portrays structure with information at each voxel (volume element), information about the interior of an object is known. This differs from traditional surface ("polygon") rendering methods and turns out to be invaluable for many analysis problems. Integrated with transparency, interactive display, and motion, volume visualization is capable of providing new understanding. It also integrates well with image processing. For many empirical data sets, the future processing paradigm will be for signal processing algorithms to extract the parameters of interest, image processing algorithms to smooth and enhance the "signal," and interactive analysis of the resulting data space using volumetric methods to gain understanding of complex structures. Primarily a methodology today and the subject of much research both to develop better and faster algorithms and to integrate it into the scientific analysis process, volume visualization is one of tomorrow's tools.

Where scientific visualization tools have been successfully implemented, attention has been paid to assuring ease of use and continuity. Software, system, and devices are made to appear transparent to the user. This requires the use of standards and pseudostandards for graphics, systems, language, and communication. Examples include PHIGS+, CGM, X Windows, Unix, and Postscript. With such tools the visualization specialist assures that hardware and software modifications, upgrades, and additions are implemented without constantly forcing scientists to interrupt their research to learn new sets of commands.

Finally, the need for visualization paradigms and formalisms stands out. Even for simple tasks such as color palette selection for an animation, there are few guidelines save experience (although the work of Donna Cox et al. at NCSA provides a good starting point). Guidelines for proceeding through the analysis process for experimental data sets are similarly lacking. In the excitement over the obvious benefits of scientific visualization, few questions have been asked about the nature of perceived information and how well the human visual system actually performs. Because visualization is a new, emerging discipline, the lack of structures is not surprising, but their development is necessary and offers significant research opportunities.

These articles represent the cutting edge of an important, emerging field. They are examples of the exciting work springing up worldwide. As the development of new scientific visualization tools makes these methodologies available to more researchers, the knowledge produced by the integration of visualization into the scientific process will further benefit mankind. Taken together, these articles provide a foundation for others to build on. ☐

Lawrence J. Rosenblum is a computer scientist in the Marine Systems Branch at the Naval Research Laboratory, where he has worked since 1977. His current research interests include visualization of oceanographic data and vision for underwater autonomous vehicles. He has also developed real-time data acquisition systems for oceanographic experiments aboard ships and aircraft.

Rosenblum received a BA in mathematics from Queens College, City University of New York, in 1964. He received an MS in 1966 and a PhD in 1971, both from the Ohio State University in mathematics. He is past chairperson of the IEEE Computer Society Technical Committee on Computer Graphics and is currently vice chairperson of the Computer Society's Technical Activities Board and co-chairperson of Compusat 89, the "Interdisciplinary World of Computing" satellite symposium. He is a member of the Computer Society, ACM SIGGraph, and the American Geophysical Union.

Readers can contact Rosenblum at Code 5170, Naval Research Laboratory, Washington, DC 20375.

Distributed Visualization at Los Alamos National Laboratory

Richard L. Phillips, Los Alamos National Laboratory

The Los Alamos National Laboratory is operated by the University of California for the US Department of Energy. Among other activities, LANL performs research and development in support of the nation's nuclear weapons programs. This requires large computer simulations of hydrodynamic phenomena. Because of the volume and complexity of resultant data, methods for visualization of information are essential. This article discusses several approaches to scientific visualization at LANL.

Even the editors of the *Visualization in Scientific Computing*[1] report would agree that scientific visualization is not a new concept; only the term is new. Figure 1 presents a brief history of visualization capability, tied primarily to the evolution of graphics hardware. Looking at the prehistory part of the diagram, many of us can recall doing what we thought was scientific visualization using crude printerplots. Later, when film recorders became available, researchers like Prueitt at Los Alamos led the way in producing sophisticated slides and animated films of scientific computations.[2] But, while the ViSC concept is not new, a recent confluence of technologies, both hardware and software, is revolutionizing the way we do scientific visualization. These developments appear on the timeline beginning in the early 1980s.

In that same time frame, much of the software for producing highly realistic rendering of scientific data began to show up as resident firmware in workstations. Then, in 1988, a series of papers on volume rendering appeared[3,4] that improved on similar methods described earlier.[5] While earlier approaches to rendering (such as Gouraud shading[6] or ray tracing[6]) required the data to have some geometric structure, volume visualization techniques can directly render scattered observations in three-space by determining the optical depth from the opacity of the volume elements along the viewing path. The resulting image, when displayed on a

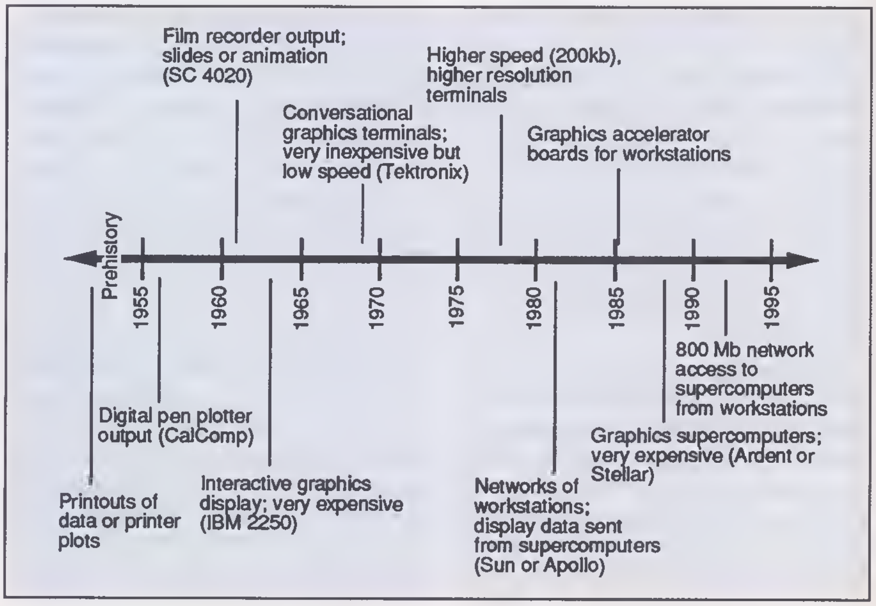
Film recorder output; slides or animation (SC 4020)
Higher speed (200kb), higher resolution terminals
Conversational graphics terminals; very inexpensive but low speed (Tektronix)
Graphics accelerator boards for workstations
Prehistory
1955
1960
1965
1970
1975
1980
1985
1990
1995
800 Mb network access to supercomputers from workstations
Digital pen plotter output (CalComp)
Graphics supercomputers; very expensive (Ardent or Stellar)
Interactive graphics display; very expensive (IBM 2250)
Networks of workstations; display data sent from supercomputers (Sun or Apollo)
Printouts of data or printer plots

Figure 1. Evolution of visualization capability.

high-resolution color monitor, shows the characteristics of the entire data set—interior and exterior—at a glance.

In the late 1980s, specialized graphics accelerators for industry standard workstations as well as graphics supercomputers were announced. As a single-user system, the graphics supercomputers are quite expensive, but they do offer unprecedented interactive visualization power. Further along the timeline, 800 megabit-per-second networks will begin to appear. The driving force in that development is the proposed ANSI standard (X3T9.3 Committee) high-speed channel (HSC) specification that Los Alamos is spearheading. Soon to become a standard, HSC hardware is already being developed by leading computer and workstation vendors. Workstations equipped with graphics accelerator boards, coupled with HSC network access to supercomputers, offer exciting opportunities for visualization, particularly in the computing environment of Los Alamos.

The Los Alamos computing environment can be characterized as widely distributed and massively interconnected. Workstations, terminals, and mainframes can all communicate with one another with relative ease. The dominant type of computing is large-scale scientific simulation using the seven Cray computers in the Central Computing Facility. To deal with the huge amount of data produced by this environment, a variety of tools and approaches have been developed for visualizing scientific information. I will describe several of these techniques, emphasizing distributed approaches to visualization.

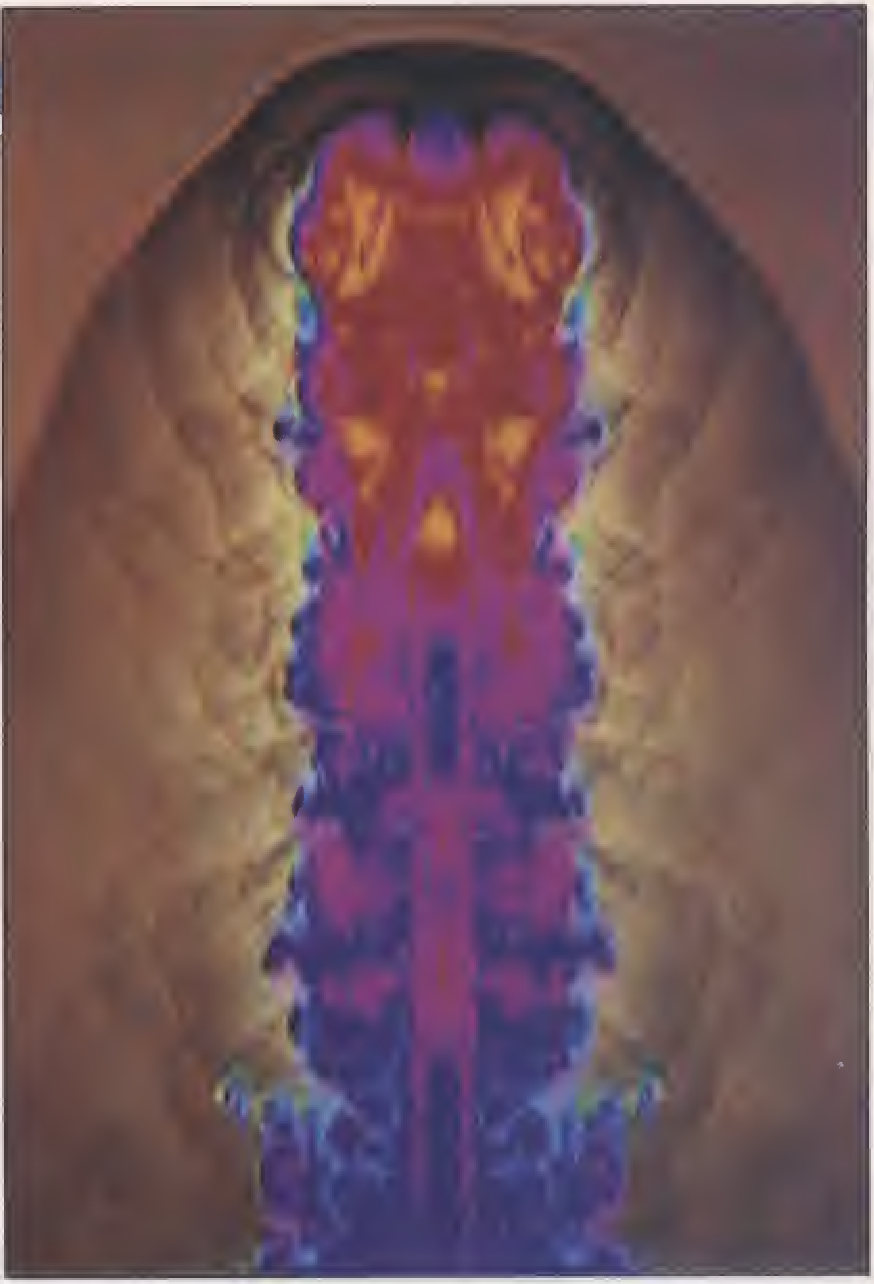

Figure 2. Numerical simulation of a gas jet penetrating a quiescent medium.

Centralized visualization: the Ultra-Speed Graphics Project

Karl-Heinz Winkler and his associates[7] have developed a video animation facility based on a Gould PN9080 system and a Gould IP8500 image-array processor. While not a distributed visualization facility, this system's capabilities have whetted the appetites of Los Alamos users for an equivalent tool. With it, complex finite difference simulations produced on a supercomputer can be viewed as 8-bit-per-pixel pseudocolor images at a resolution of $1,024 \times 1,024$ at 8 frames per second or 512×512 at 30 frames per second. Arbitrary color lookup tables can be specified to induce desired features to emerge. This facility forms the basis for Winkler's Los Alamos Numerical Laboratory, which is used not only by Los Alamos researchers but by scientists from throughout the world. The facility allows investigators to explore huge collections of simulation data in a relatively short period and occasionally has led to the discovery of unexpected phenomena.

The system is organized so that data sets computed on the Crays are transferred to the display system over a 27-megabit-per-second parallel interface. While this data rate is high enough to allow continuous monitoring of typical numerical experiments, the image support processor is used primarily to receive large sets of compressed solution numbers from several Crays, turn them into individual images at a rate of up to 100,000 images per day, and display them as movies on digital display devices. Storing the images in digital form completely decoupled from a particular color representation allows tremendous flexibility in digging out features hidden in the numbers. Figure 2 graphically shows

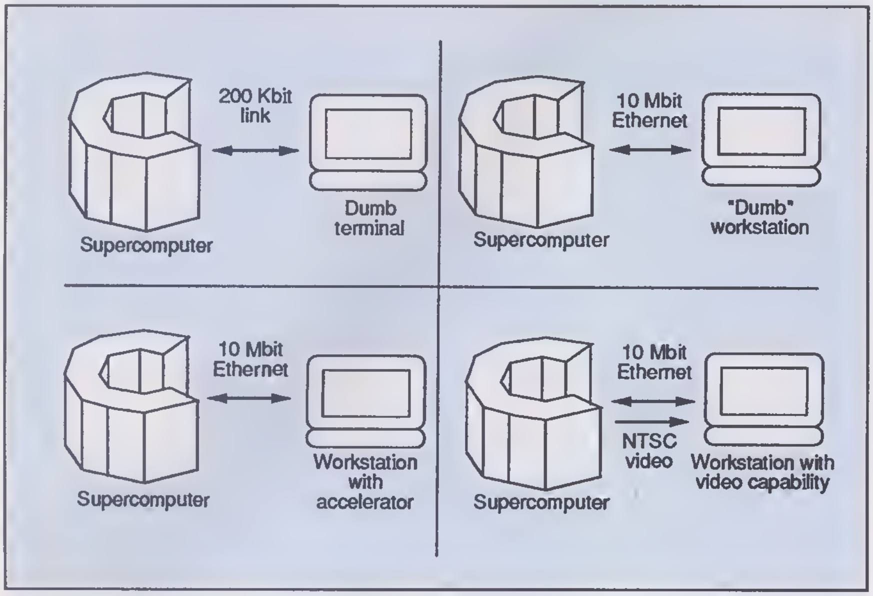

Figure 3. Supercomputer/display connections.

the computer simulation of a high-speed jet entering a quiescent ambient gas.

While the Winkler work is ground breaking and important, it represents an approach to visualization that only a few can afford; such a facility is clearly expensive to replicate. Other approaches to visualization at Los Alamos seek to bring the visualization environment into the scientist's office, using relatively low-cost workstations and network connections to the Crays. There are several ways to accomplish such a connection; some are discussed below. Many of these approaches have been inspired by, or are an outgrowth of, the Ultra-Speed Graphics Project.

General approaches to distributed visualization

Figure 3 depicts several approaches to connecting graphics displays to supercomputers. They are all in use or under development at Los Alamos.

The connection shown in the upper left is typical of one still in use at Los Alamos. In particular, the terminal would be one of the Tektronix family, a 4014 or 4125. A special high-speed interface provides 300-kilobit-per-second access to a Cray and, because of the efficient data format used by Tektronix terminals, complex images can be produced quickly. There is no opportunity, however, to perform any local processing with these systems. All data is produced and sent by a supercomputer.

Simple workstations connected to a supercomputer (upper right diagram in Figure 4) offer many advantages over a simple terminal. Since there is local processing power, the user can save data sets produced by the supercomputer and manipulate them further off-line. However, the graphics performance of many low-end workstations is inferior to the terminal arrangement discussed above. Although the local area network connection (Ethernet) to the Cray has a bandwidth of 10 megabits per second, the actual throughput is much less, typically 0.5-1 megabit per second. This is partly due to protocol overhead, but also because what is referred to in the figure as a "dumb" workstation typically has poor drawing performance. Examples of these are Motorola 68000-based workstations with no graphics acceleration hardware.

But there are workstations with superior graphics performance; the silicon-enhanced systems from Silicon Graphics are but one example. You can also equip a workstation with a graphics accelerator board. Workstation manufacturers provide these, as do third-party hardware suppliers. Recently, Alliant Computer Systems and Trancept Systems (now a subsidiary of Sun Microsystems) introduced accelerators that promise to have a strong impact on visualization efforts. Alliant offers the GX4000, a board set that plugs in to a standard Sun Microsystems workstation. This board implements the proposed PHIGS+ standard in microcode and thus provides a 3D graphics capability for the host workstation. This means that full 3D data can be sent from a supercomputer and transformed locally. Thus, an object can be animated just by transforming it and redisplaying it in a new orientation. Depending on the complexity of the object, good to excellent animation rates can be achieved, that is, 4 to 24 frames per second.

The Trancept board is called the TAAC-1, an acronym for Trancept Application Accelerator. As the name suggests, this board can be used to speed up all kinds of local computing tasks. When used for graphics applications, it can perform sophisticated rendering operations at very high speeds. While the board is an add-on to a Sun Microsystems workstation, it has its own frame buffer and is capable of producing high-resolution 24-bit-per-pixel images. The board also has the potential of accepting data directly (bypassing the Sun processor) at 600 megabytes per second. Work under way at Los Alamos to exploit this feature will be described later.

While the graphics performance of both of these accelerators outstrips conventional workstation hardware, the data delivery rate to them is, in general, still limited by Ethernet throughput. This may not be a problem for investigations where a relatively small amount of object data is sent from the supercomputer and is expanded locally to produce high-quality images, animated sequences, or both.

There is a large class of problems, however, where this is not an especially useful approach. These are typically simulations involving finite difference calculations where complex meshes are used to describe the spatial aspects of the problem. An example of such a mesh is shown in Figure 4.

This represents one instant in time in a sequence of many thousands of such images. As the problem evolves, some part of the mesh changes from one time step to another. Each image comprises several thousand data points and, due to Ethernet limitations, you cannot expect to achieve very good graphics display rates. Accelerator boards are no help here because there is no local algorithmic processing they can usefully perform. An exception would be the case where only those parts of the image that change in a time interval are sent to the workstation. The accelerator would maintain the graphics database and apply the changes locally. This is not a simple task, however, and has not been done in practice.

Another approach to remote data visualization is to perform some, if not all, of the graphics processing on the supercomputer and send the complete representation to the workstation as an image, that is, pixels. But, if the image resolution is say, 512 × 512 × 8, one-quarter megabyte of data must be sent for each display. Such a volume of data would seem to preclude this approach, at least in the Ethernet environment described above. There are ways to circumvent the brute force approach, however. One is to buffer several images in the workstation's virtual memory and then enter a local animation loop. The other is to convert the images to video at the supercomputer and provide an alternate display link to the workstation. Both approaches are being pursued at Los Alamos.

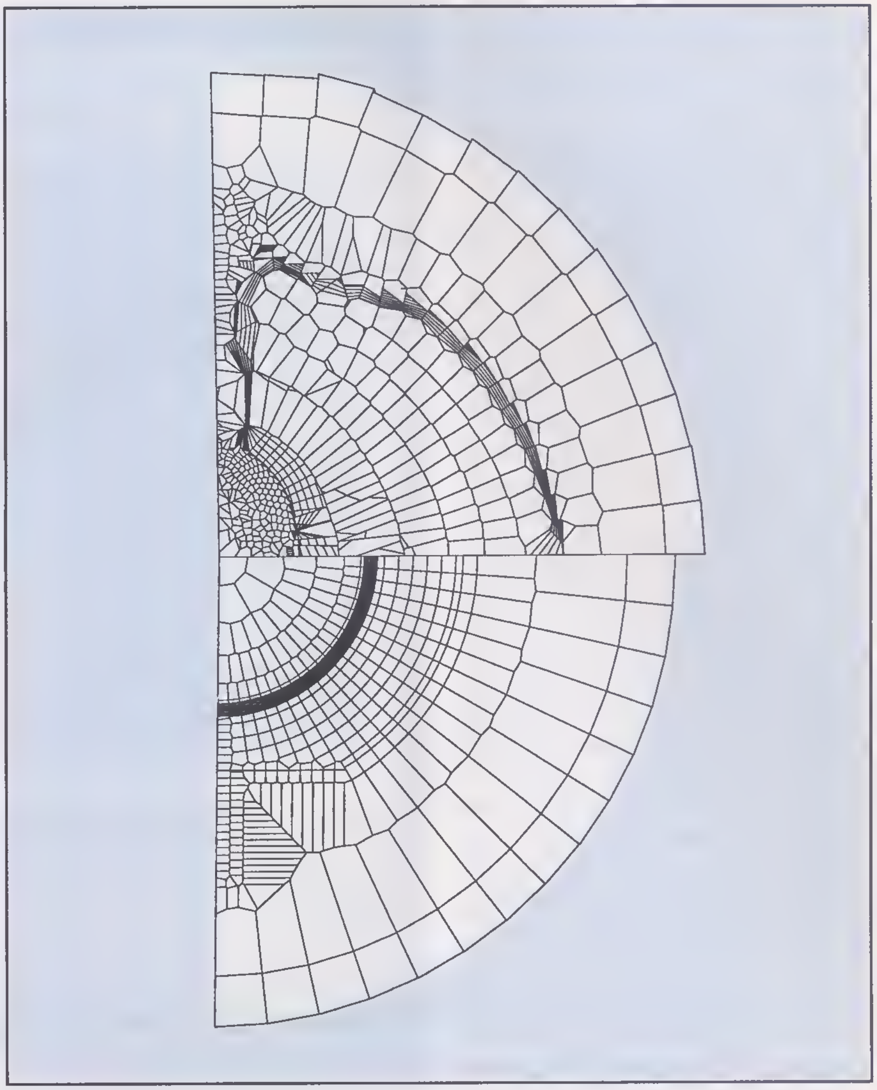

Figure 4. Complex finite difference mesh.

Specific distributed visualization projects

This section describes two distributed visualization projects—Simple Interface and the Scientific Visualization Workbench—under way at LANL and discusses implications for the future.

SI: An interface between a supercomputer and a workstation. Saltzman[8] has developed a system for viewing pixel images produced on a supercomputer on a standard Sun workstation; no special hardware is required. SunView, the standard Sun window system, runs on the workstation. Communication with a Cray is via Ethernet.

SI (for Simple Interface) is based on a client/server model. The Cray runs an SI server that listens for requests from the client side running on a Sun. When the Sun requests image data, the server opens a communication path (via sockets) and sends the requested information. Once done, the server breaks the path and goes about listening for additional requests. At this point, the program on the workstation can peruse the frames of data it has received, either a frame at a time or in an animation loop. Figure 5 shows SI in use.

SI has the advantage of requiring only standard hardware and software and, thus, can be widely used. At Los Alamos, the nearly 700 Sun workstations make the ubiquity of use important. SI allows only 128 colors to appear in an image, however, and only 10 frames may be cached locally. If these limitations are not significant, SI represents a very effective distributed visualization tool.

Visualization using video images. The lower right panel in Figure 4 shows a workstation connected to a supercomputer by both Ethernet and a National Television System Committee video link. Of course, the workstation must have a device that can handle video data. This is a type of accelerator, a board that accepts an image directly and displays it on the workstation screen. Parallax Graphics makes such a board. It can digitize an NTSC video signal in real time, that is, 30 frames per second. The Scientific Visualization Workbench[9] developed at Los Alamos makes use of the Parallax board capabilities, in conjunction with an already-developed supercomputer frame buffer system[10] and the News window system.

Fowler and McGowen developed the supercomputer frame buffer system. The frame buffer attaches directly to a Cray channel capable of delivering data at 48 megabits per second. The developers visualized the system being used by communicating with the Cray via a dumb terminal at 9,600 bits per second. Supporting software then converts graphics metafiles to raster images and sends them to the frame buffer channel. Images are viewed on an RGB (red, green, and blue) monitor at up to 26 frames per second. The RGB signal can be displayed remotely by extending it to a user's office by fiber optics. The Scientific Visualization Workbench takes this one step further by bringing the video signal into a window on a workstation screen.

This image-based approach to visualization overcomes the speed limitations of other approaches, but at a price. Image quality is lost by rasterization for the frame buffer and certainly by NTSC encoding of the video signal. Thus, you trade image fidelity for bandwidth.

On the other hand, the user can see images as fast as they can be produced by the supercomputer—as high as 26 frames per second—and temporal fidelity is often more important than image quality. The flexibility to experiment with new visualization algorithms, rather than relying solely on rendering methods provided by the workstation vendor, is another advantage to be gained from processing graphics data on a supercomputer.

The structure of the Scientific Visualization Workbench is depicted in Figure 6.

The workbench comprises a Sun Microsystems 3/160C workstation equipped with a Parallax 1280 Series Videographics Processor, a Sony high-resolution color monitor, a Cray Research supercomputer equipped with a solid-state drive (SSD), and a 512h × 512v × 8 bits-per-pixel framebuffer attached to the Cray by a 48-megabit-per-second channel. Communication between the Sun and the Cray is via Ethernet using TCP/IP protocol. The Cray runs Unicos 3.0 (System V-based Unix), and a special Parallax-enhanced version of the News 1.1 window system (Pnews) runs on the Sun. A News client program on the Cray controls the transmission of rasterized data to the frame buffer and ultimately to the workstation over an NTSC video link. Because of the Postscript underpinnings of News, it is possible to produce color or black-and-white hard copy of the video images of anything on the screen. The magnified workstation screen in the lower left of the figure suggests the presence of several types of windows, in which independent

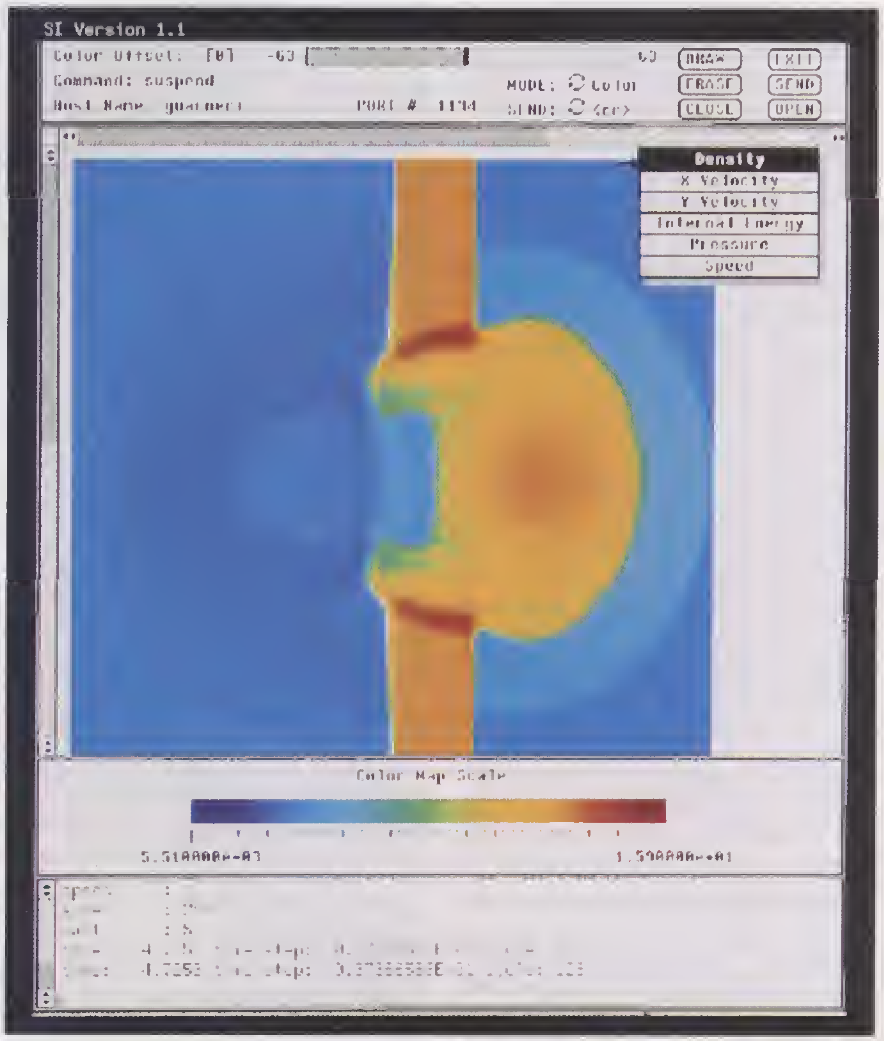

```
SI Version 1.1
Color Offset:  [0]   -63                    63
Command: suspend
Host Name  guarneri          PORT #  1194
MODE:  Color
SEND:  <cr>
```

Figure 5. Simple Interface in use.

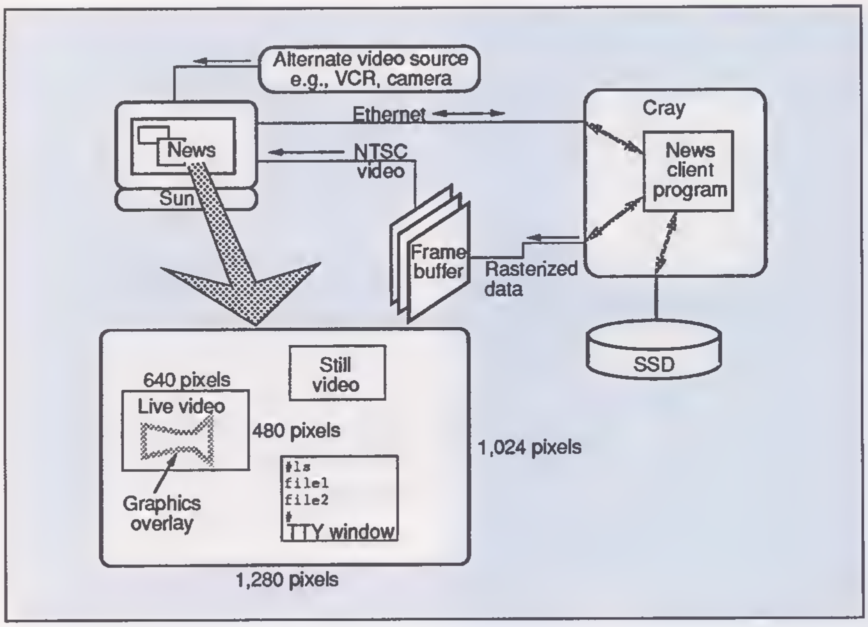

```
#ls
file1
file2
#
```

Figure 6. Scientific Visualization Workbench.

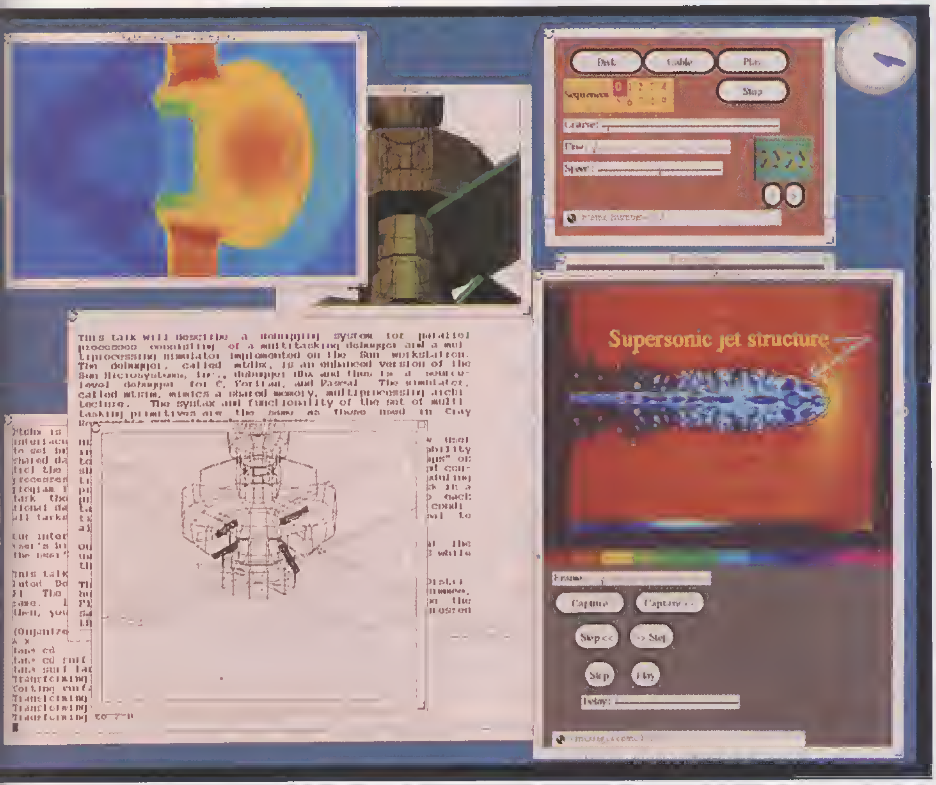

Figure 7. Screen dump of Scientific Visualization Workbench display.

processes run simultaneously. Figure 7 is an actual screen dump of the workbench in use.

Higher speed networks: Implications for the future. Los Alamos is leading the specification and development of an 800-megabit-per-second channel, the HSC (high-speed channel)[11] for use in a fiber-based, point-to-point crossbar network. DEC, IBM, Sun, Tektronix, and others have already committed to building HSC hardware, while DEC is playing a key role in the development of the network, called CP*. Figure 8 shows a diagram of the future Los Alamos Integrated Computer Network (ICN), based upon the HSC and CP* technology.

With the availability of such high network bandwidth, many distributed visualization schemes become possible. For example, one need no longer be content with NTSC video or require a separate video link. CP* will easily accommodate CCIR 601 digital video or even the emerging high-definition television (HDTV) standard of 1,100 lines. The Scientific Visualization Workbench would be upgraded to account for that capability.

Another approach being explored involves a Sun workstation equipped with a Sun TAAC board modified to operate at HSC speeds. With this system, you can transfer full floating-point, 3D graphics data sets from the supercomputers at very high speed and use the local processing power of the TAAC board for fast rendering of the data. This configuration is being tested over a direct fiber link between a workstation and an IBM 3090/200 equipped with an HSC interface.

Many changes are under way within the Los Alamos computing infrastructure to improve capabilities for distributed visualization. As the CP* network is put in place, gateways to Ethernet and FDDI (Fiber Distribution Data Interface) fiber networks will also be provided. The CP* backbone will provide the bandwidth for transmission of 24-bit color images with a resolution of 1,024 × 1,024 at a rate close to 30 frames per second. This will permit local office viewing of images generated on a supercomputer, using a high-resolution

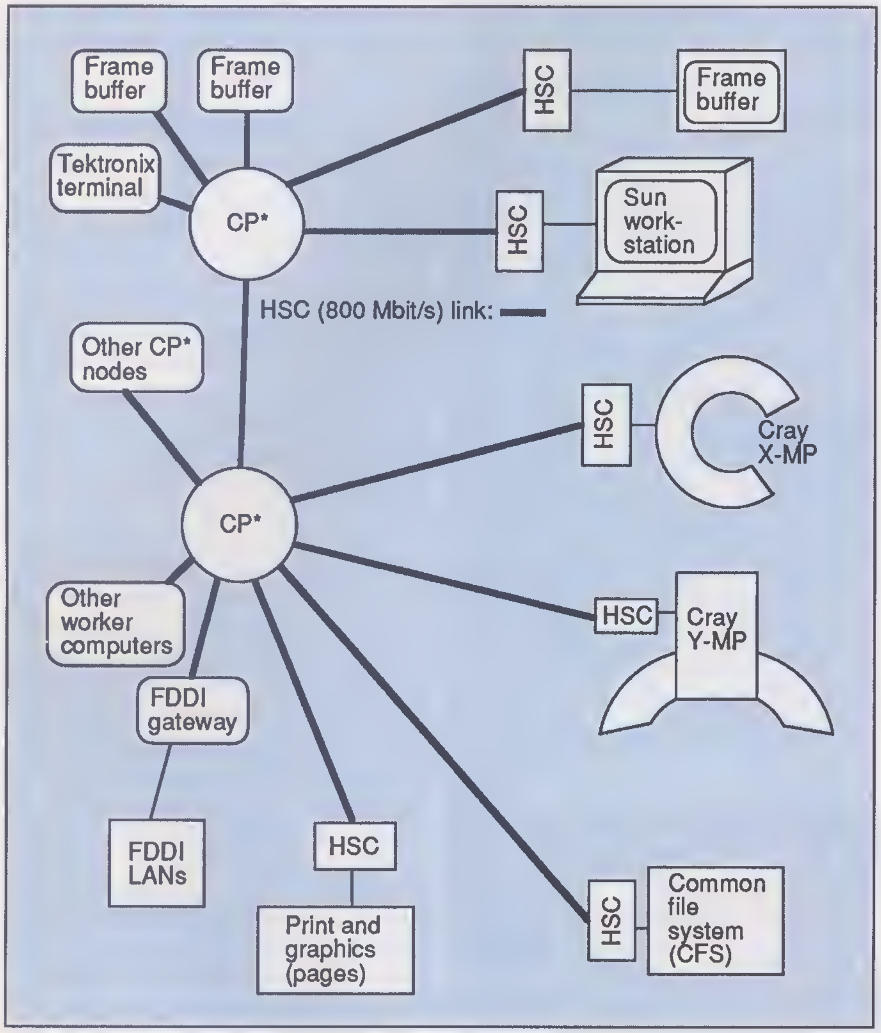

Figure 8. Future Integrated Computer Network with high-speed channel and CP*.

successor to the Scientific Visualization Workbench. Alternatively, users will be able to transfer data sets of 3D, floating-point information at a rate of 100 megabytes per second to accelerator-equipped workstations or to 3D workstations like the Silicon Graphics Iris for local rendering.

As in many organizations, security is an important issue, and the attendant necessary firewalls are often a barrier to effective interactive use of computing facilities. At Los Alamos, this issue will be resolved in early 1990 by the introduction of a secure IP (Internetwork Protocol) router. Then, users in all parts of the computing environment will enjoy high-speed distributed visualization.

Exploratory work on the use of massively parallel computer architectures for scientific visualization is just beginning. The recent installation of a Thinking Machines CM2 with 65,536 processors will accelerate this effort. One of the CM2 frame buffers has already been installed on the cable video network mentioned earlier, making its output available for use with the Scientific Visualization Workbench.

Finally, most of the techniques and systems described here are portable to other installations. The Scientific Visualization Workbench can be easily duplicated or implemented on different hardware, using different software. The Simple Interface can be installed wherever SunView is available on workstations and RPC (remote procedure call) is supported on the network. Even HSC, being an ANSI standard, and CP*, which will be publicly available, can be implemented elsewhere.□

References

1. "Visualization in Scientific Computing," *Computer Graphics,* B.H. McCormick et al., eds., Vol. 21, No. 6, Nov. 1987.
2. M.L. Prueitt, "A Window on Science," *IEEE Computer Graphics and Applications,* Vol. 7, No. 9, Sept. 1987, pp. 4-8.
3. C. Upson and M. Keeler, "V-Buffer: Visible Volume Rendering," *Computer Graphics,* Vol. 22, No. 4, Aug. 1988, pp. 59-64.
4. R.A. Drebin, L. Carpenter, and P. Hanrahan, "Volume Rendering," *Computer Graphics,* Vol. 22, No. 4, Aug. 1988, pp. 65-74.
5. G.T. Herman and J.K. Udupa, "Display of 3D Digital Images: Computational Foundations and Medical Applications," *IEEE Computer Graphics and Applications,* Vol. 3, No. 5, Aug. 1983, pp. 39-46.
6. J.D. Foley and A. Van Dam, *Fundamentals of Interactive Computer Graphics,* Addison-Wesley, Reading, Mass., 1982. See Chapter 16.
7. K.-H. Winkler et al., "A Numerical Laboratory," *Physics Today,* Vol. 40, No. 10, Oct., 1987, pp. 28-37.

8. J. Saltzman, "Simple Interface (SI): An Interface Between a Cray Supercomputer and Sun Workstation," LA-UR-88-2171, Los Alamos National Laboratory, June 1988.
9. R.L. Phillips, "A Scientific Visualization Workbench," *Proc. Supercomputing 88,* Nov. 14-18, 1988, Orlando, Fla., pp. 148-155, CS Press, Los Alamitos, Calif., Order No. 882.
10. J.D. Fowler and M. McGowen, "Design and Implementation of a Supercomputer Frame Buffer System," *Proc. Supercomputing 88,* Nov. 14-18, 1988, Orlando, Fla., pp. 140-147, CS Press, Los Alamitos, Calif., Order No. 882.
11. B. Morris and D. Tolmie, "High-Speed Channel (HSC): Mechanical, Electrical, and Signalling Protocol Requirements," *Draft Proposed American National Standard, X3T9.3/88-023,* Aug. 1988.

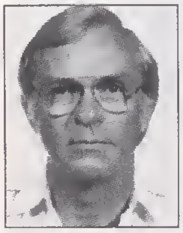

Richard L. Phillips is a staff member in the Computer Graphics Group of the Computing and Communications Division of the Los Alamos National Laboratory. His current research interests are scientific visualization, multimedia workstations, distributed computing, window systems, and electronic publication.

Phillips received a BSE in mathematics in 1956, an MSE in aerospace engineering in 1957, and a PhD in aerospace engineering in 1964, all from the University of Michigan. He is treasurer of ACM SIGGraph and played major roles in the 1978, 1986, 1987, and 1988 SIGGraph conferences. He is a member of the IEEE, the IEEE Computer Society, and ACM.

Phillips can be reached at Los Alamos National Laboratory, PO Box 1663, MS B-272, Los Alamos, NM 87545.

Visualization Tools at Lawrence Livermore National Laboratory

Brian Cabral and Carol L. Hunter, Lawrence Livermore National Laboratory

Lawrence Livermore National Laboratory (LLNL), operated by the University of California for the US Department of Energy, supports research programs in nuclear weapons, magnetic fusion energy, laser fusion, laser isotope separation, biomedical and environmental sciences, and applied energy technology. These programs require research in basic scientific disciplines, including chemistry and materials sciences, computer science and technology, engineering, and physics.

Scientists have at their disposal seemingly limitless data and computing power, and they are always seeking enhancements of their ability to access, explore, represent, and manipulate data. Scientific visualization is an emerging interdisciplinary field that is responding to the need for insight into data representations through the creation of visualization tools.[1] There is no "right" way to provide insight, and the user directly benefits from different efforts at scientific visualization. In this article, we describe two efforts to develop visualization tools at Livermore, one within the Computation and Defense Sciences Departments and the other within the Engineering Department.

Graphics in the supercomputer environment

For 40 years, scientists have produced numerical simulations with computers and derived insight from calculations. During this time, increases in computing capability have accompanied increases in calculation complexity and computer graphics sophistication. However, many graphics methods from the 1950s are still adequate for decision-making today and are vital tools to the scientist. Visualization tools range from *x,y* plots to animated three-dimensional renderings. Visualization tools should include any and all data representations and allow a researcher to choose between a simple plot and a complex image based on deadlines and detail constraints.

In the quest for faster, more powerful, and more responsive computation, Livermore developed a complex supercomputer environment over the years, managed by the Livermore Computer Center (LCC). The environment includes seven Crays, hundreds of workstations, and the interconnecting network structures. Graphics for the 2,000 supercomputer users is provided by the Television Monitor Display System, an 18-year old DeAnza central frame buffer system. TMDS provides black and white, low-resolution (512 × 512) service throughout the one-mile-square facility, at up to 30 frames per second. It is not feasible to upgrade to high resolution and color or extend the network with the existing hardware. In 1988, an LCC-chartered committee made recommendations for the functionality required to replace the TMDS.[2] We will describe our prototype efforts to provide information on the performance of the TMDS, using high-resolution color systems that add functionality and interactivity.

Most LLNL research involves considerable computational effort. Massive amounts of data are consulted repeatedly in various transformations and representations. This work involves frequent comparisons—comparing new calculations to old calculations, to previous time steps, to experiments, to calculations on different software, and to analytic solutions. An essential component of scientific visualization is on-line browsing into multiple windows to allow comparison of data and the ability to follow a train of thought on the screen. Scientific visualization must address quantitative as well as qualitative information. At any point in browsing through graphics images, the scientist needs to be able to query the graphics system as to the specific values of data being displayed. Many different easy-to-use tools must be available to improve the scientist's everyday, analysis graphics capabilities.

Magic. Magic, or Machine Graphics in Color, is designed to bring high-resolution, color images to the offices of researchers over a high-speed network. The Magic network (Figure 1) transfers Cray graphics data up to one kilometer at a throughput rate of 2.5 megabits per second. When we started the project, we wanted to deliver Cray supercomputer graphics to advanced graphics terminals in researchers' offices at a high speed using low-cost, off-the-shelf products.[3] We chose Tektronix 412x graphics terminals because they offered high plotting speeds, high resolution, a wide range of color, picture segmentation, flexible graphics-input tools (cursor, mouse), fill-area, raster-plotting functions, and three-dimensional capabilities. The terminals were a known and reliable medium. The network solution was an unknown.

Magic provides a platform for high-level data representation and algorithm experimentation. Users understand their problem domain; our local graphics experts understand graphics. The experts developed a variety of data representations and worked with the users to identify the most informative representations. Interest was most often sparked when the graphics showed an unusual feature that looked like a "bug in the graphics," but 95 percent of the time they turned out to be a bug in the user's application. Variations of color tables proved to be extremely informative; the use of a rainbow color table and a gray-tone color table were most productive. Color presents the opportunity of displaying multiple parameters of the user's data (Figure 2).

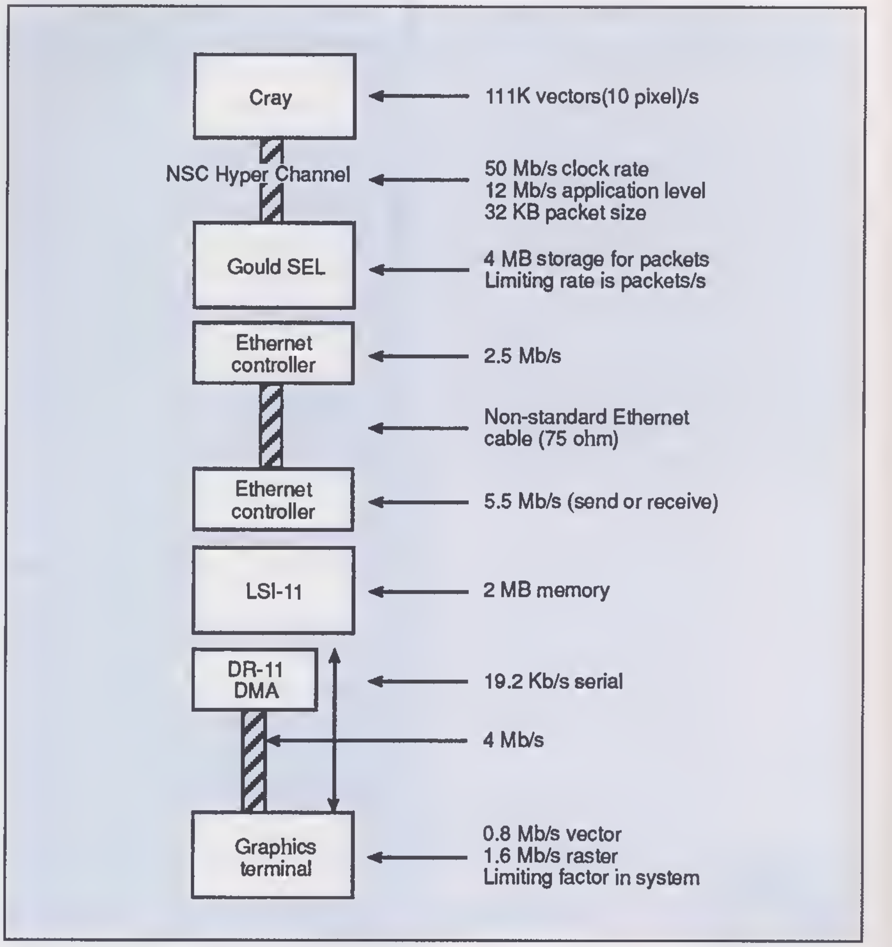

Figure 1. The Magic network transports data from any Cray in the LCC, over HyperChannel to a Gould SEL computer being used as a graphics concentrator, over a combination of existing 75 ohm coax and 50 ohm coax to remote buildings where an LSI-11 buffers the data to a Tektronix graphics terminal.

A follow-on effort with the Magic network has provided real-time animation with the use of a Super VHS laser video disc recorder. Frames are captured from the Tektronix and replayed on the video disc player. A second follow-on effort, MacMagic, involves sending raster images from the Cray environment to a Mac II over the Magic network. ImageTool,[4] developed by the National Center for Supercomputing Applications, proved to be an extremely powerful public-domain product; this package provides the ability to manipulate the palette, view multiple images, and animate sets of image frames. ImageTool was modified to accept a message stream from the Magic network and display raster images sent from the Cray.

The Graphics Workstation Project. The Graphics Workstation (GWS) Project moves the user's graphics environment from the supercomputer to a workstation and provides interactivity, a friendly user interface, and additional functionality.[5] In our initial project, we used four Sun-3/160 graphics workstations with floating-point and graphics accelerators and a color and monochrome monitor for each workstation.

The most important tool developed to date is a postprocessing program for two-dimensional data on a rectangular mesh. The screen displays panels with push buttons, sliders, pop-up windows, and a zoom and pan feature. Typical plots are line and pseudocolor plots of various physical data (temperature, pressure, density, etc.), mesh plots, velocity-vector plots, and material-boundary plots, all of which can be superimposed.

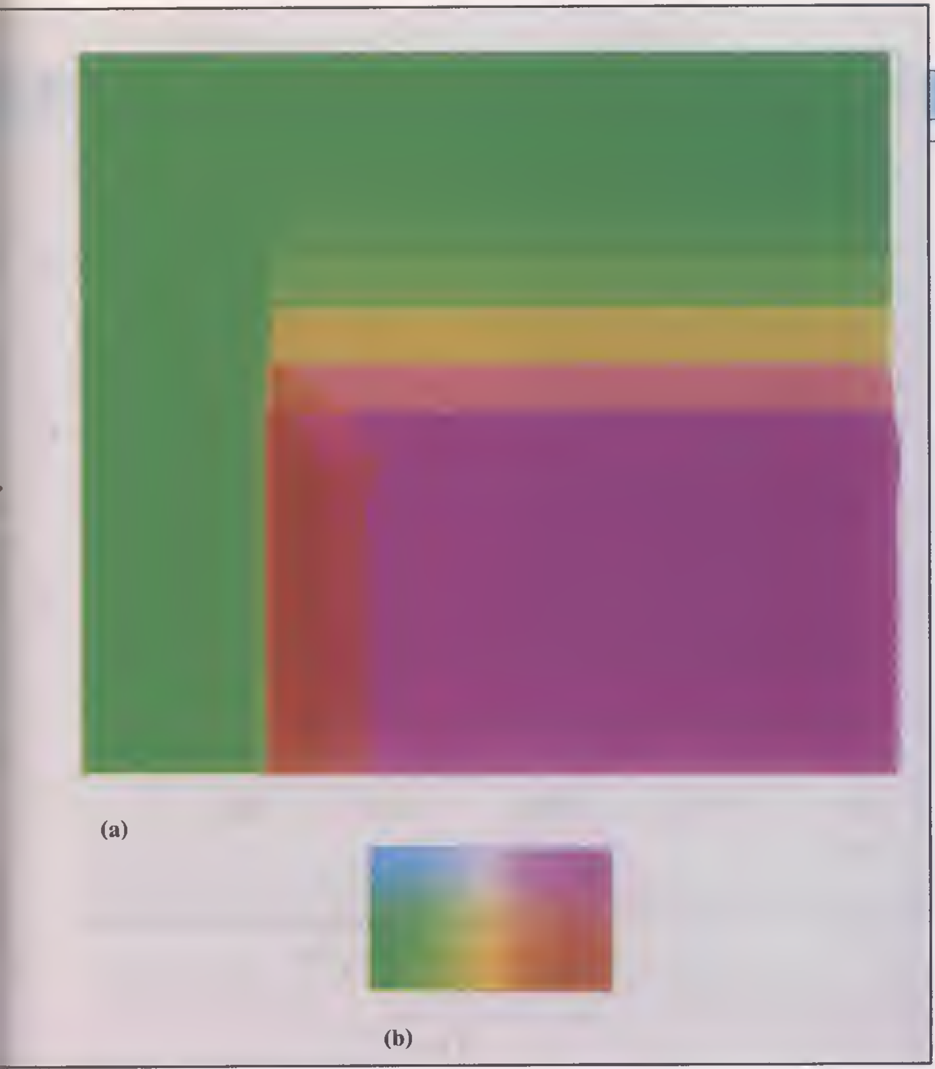

Figure 2. Two different physics variables are represented on a two-dimensional grid (a) by using a two-dimensional color table (b), where the horizontal axis of the color table represents one variable, and the vertical axis represents the other variable. The upper right corner of the color table represents high values of both physics variables.

As many as four different frames can be displayed at one time (see Figure 3) for comparison and follow-train-of-thought capabilities. A pick-and-query capability (Figure 4 and 5) is designed to determine the exact value of coordinates and physical data (quantitative information) at any point on the screen. Other GWS tools include a metafile viewer to display previously stored graphics images; a three-dimensional surface plotting routine (Figure 6); and a Cray-to-Sun system that allows a Cray application to send graphics from a Cray graphics library over the network to a remote Sun workstation window.

The GWS project is a success because it demonstrates the feasibility of workstation graphics for supercomputer data and provides essential capabilities—comparison, train-of-thought capabilities, and both qualitative and quantitative information. However, the performance of GWS was inadequate: a faster CPU, more memory, and higher performance graphics were required. The main goal of the current phase of GWS is to obtain high-performance graphics from portable graphics software based on standards.[6] In 1988, the GWS project ordered two Stellar GS1000 graphics workstations because they integrated a powerful floating-point capability with very high speed graphics. Our high-performance graphics and workstation tools will be based on a standard software platform that includes Unix, C, PHIGS, and X Windows.

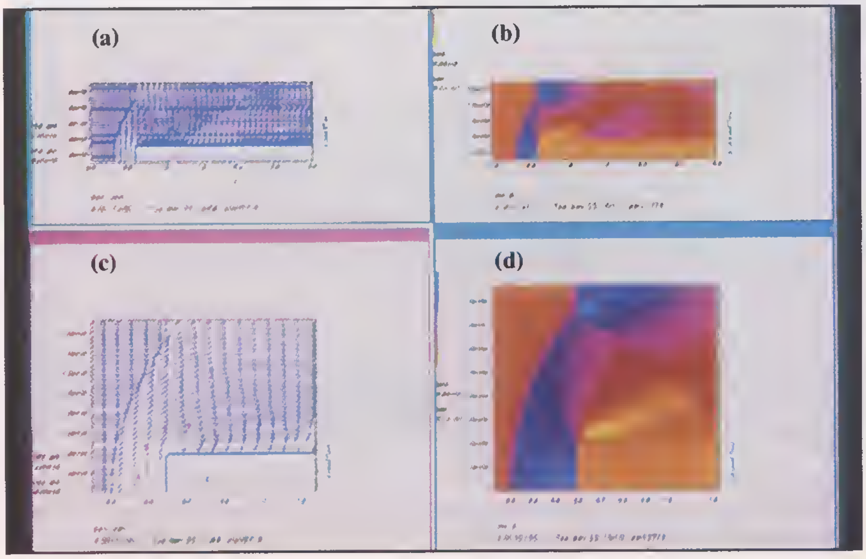

Figure 3. Multiple images display a wind tunnel with different data representations, allowing the user to follow a train of thought and compare images; (b) is a zoomed image of the (a) material, contour, and vector plot, and (d) is a zoomed image from the (c) pseudocolor plot.

The Advanced Visualization Research Project

The Advanced Visualization Research Project at LLNL addresses changing technological constraints, while allowing for rapid research, development, and the delivery of new visualization algorithms to engineers and scientists. AVRP is currently involved with ongoing work in the areas of parallel graphics algorithms and visualization system integration. Through vigorous research and the development of new visualization algorithms, AVRP strives to influence the direction of scientific visualization.

In this section, we describe an existing software system that allows scientists to experiment with state-of-the-art visualization algorithms. The considerations that drive the system's architecture include the varied data structures used in computational models; the need for interactivity; the speed, capabilities, and diversity of the laboratory's graphics hardware; and the need to rapidly prototype new visualization algorithms.

A function-based taxonomy. We are developing new methods to transform symbolic and numerical data into a visual form. By abstracting the symbolic and numerical data into a general taxonomy, we define an interface between scientific applications and visualization tools. Much of the analysis done by scientists compares results of experiments and/or simulations. These results can be described as a function of *m*-independent variables.

Usually, an independent variable is one of three spatial variables, *x*, *y*, and *z* and the temporal variable, *t*. Our taxonomy characterizes scientific data as a set of functions with an arbitrary number of independent variables. It further characterizes these functions based on the geometric spacing (Eulerian geometries are a rectangular fixed lattice; Lagrangian geometries are an arbitrarily spaced moveable lattice) of the independent variables.

The AVRP software architecture is composed of four major components (see Figure 7). At the user level, a set of tools and algorithms graphically render scientific data sets. The user interface and

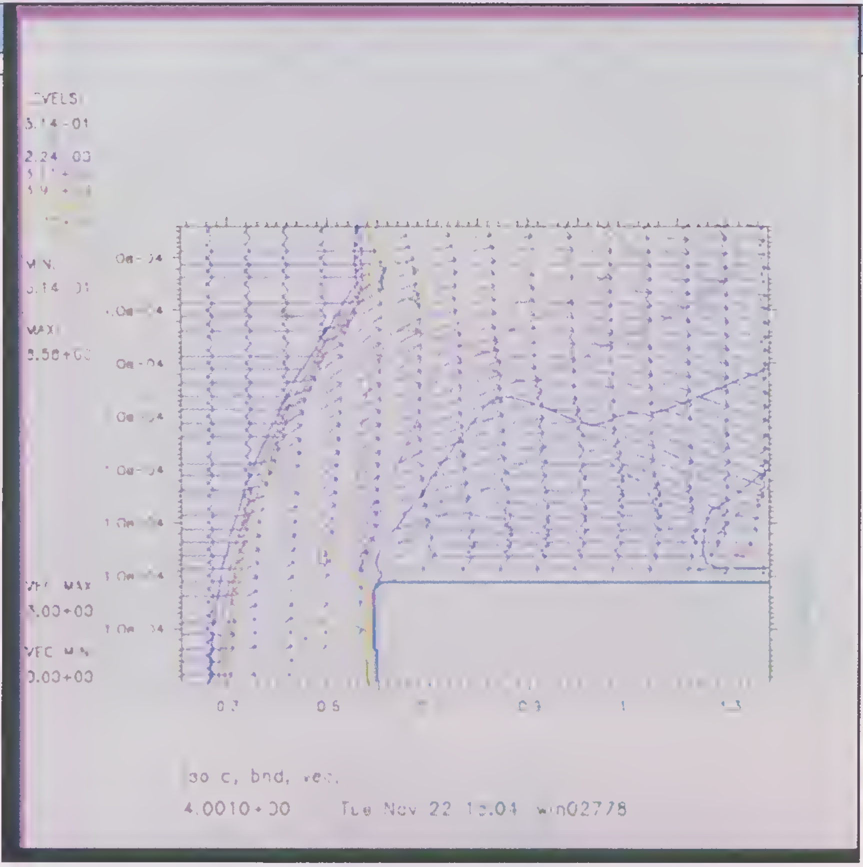

Figure 4. View of the wind tunnel calculation in which velocity vectors are superimposed over material interfaces and density contours. B, C, D, E, F are pick-and-query characters, which relate to the text window in Figure 5.

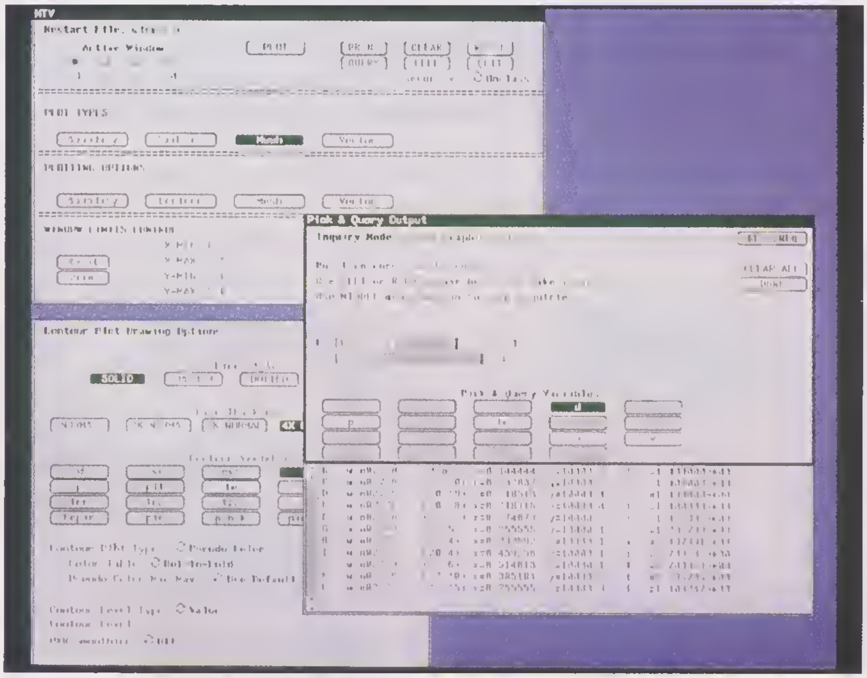

Figure 5. Screen dump of the monochrome screen showing control panel, contour selection panel, and pick-and-query results of the user's quantitative search of the data.

graphics device abstraction are provided for all AVRP tools by Gab (Graphics Abstraction) and Glib (Graphics Library). Cell IO provides the scientific data abstraction that allows us to build a general interface to the AVRP tools.

By characterizing scientific data with the described taxonomy, we have decoupled scientific applications from specific visualization techniques. This abstraction allows us to build a general interface to many scientific applications. The Cell IO software supports a file format that accepts multiple functions represented as two-dimensional and three-dimensional Eulerian grids of data. The engineer may specify an arbitrary number of dependent variables per cell, multiple time steps, and multiple grids (that is, functions) of any discrete resolution. Dependent variables may be used in any combination within a single grid, and the number and type of dependent variables may vary from grid to grid.

The Cell IO data abstraction defines the structure of the data and the syntax of the file, rather than the semantics of the grid data. Thus, it is possible for one tool to expect a two-dimensional grid and another to expect a three-dimensional grid. Most tool/application interface problems occur because of the nonuniformity of syntax and data structures required by various tools or subroutines. This software makes it easier for an application to interface to our visualization tools than with methods based on graphics libraries. Our architecture provides the scientific programmer a high-level, easy-to-use, generalized interface to visualization tools.

Glib is a three-dimensional graphics package based on a low-level graphics system (such as X Windows or Silicon Graphics GL graphics package). Our philosophy is to provide, via Glib, the graphics functionality needed by our tools and to avoid the excess software of the high-level graphics libraries. Gab is a set of high-level windowing abstractions built upon Glib. By basing all our tools on Gab and Glib, we can guarantee the performance and behavior characteristics of the underlying system. As the X Windows user interface tool kits and three-dimensional extensions mature, the need for Gab and Glib may disappear.

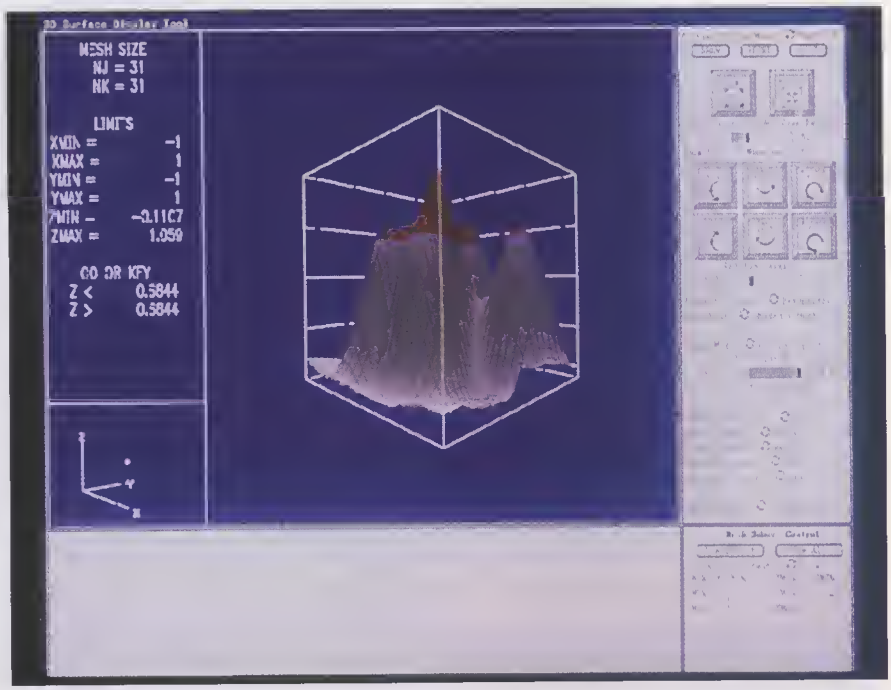

```
MESH SIZE
 NJ = 31
 NK = 31

 LIMITS
XMIN =        -1
XMAX =         1
YMIN =        -1
YMAX =         1
ZMIN =   -0.1107
ZMAX =     1.059

COLOR KEY
 Z <      0.5844
 Z >      0.5844
```

Figure 6. An alternative tool provides 3D rendering of the wind tunnel calculation. The tool provides quantitative as well as qualitative information via the bicolor split feature, allowing the user to select the *z*-value to be used for the split.

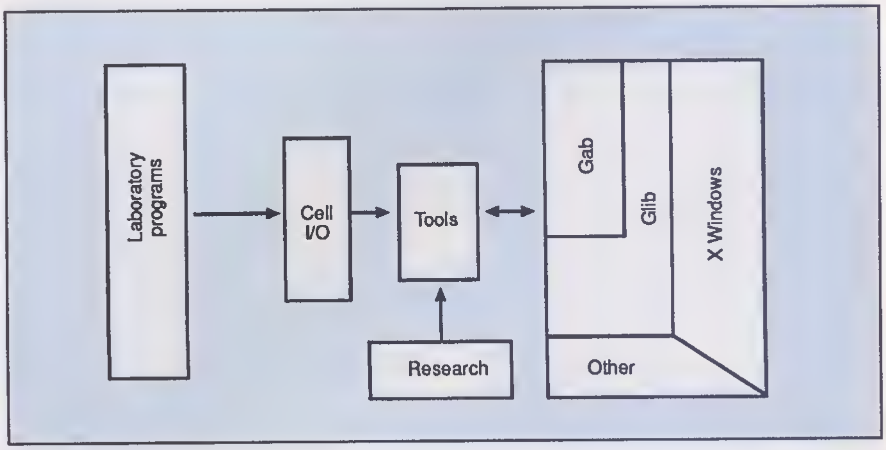

Figure 7. The AVRP software architecture.

Tools and algorithms. The purpose of the AVRP tools is to provide our scientists with new visualization algorithms. All of the tools are based on the Cell IO, Gab, and Glib software detailed above. Next, we describe four of the tools from the AVRP system. Our tools can display two-dimensional time sequences, time sequences of contours of three-dimensional grids, and various components of three-dimensional finite difference meshes, and can volumetrically render three-dimensional grids. Most tools provide interactive rotation, zooming,

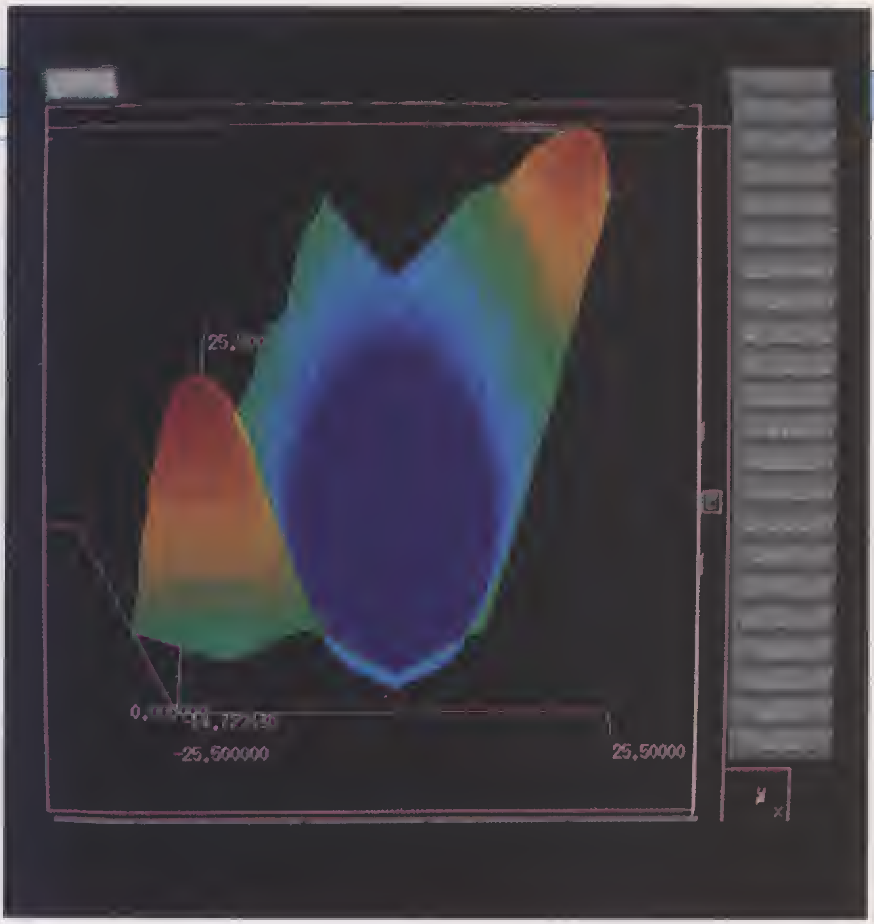

Figure 8. An interactive Surface display of a probability field of quantum particles represents the area of the box covered by the crystal.

panning, and other geometric manipulations.

We can interactively display two-dimensional meshes of data with an arbitrary number of dependent variables using Surface. Any two of these dependent variables can be displayed at one time, one spatially and the other with color. Since any two independent variables define a mesh, we can display any tensor product surface. Users interactively choose the dependent variables for display and select orientations at specific times in the data set. Using this tool, it is simple to make an animation of a time sequence of two-dimensional surface data (see Figure 8).

Interactive displays of contour surfaces carved from or constructed out of a three-dimensional volume are possible with an isosurface tool. The isosurface algorithm is based on the marching cubes algorithm.[7] We support multiple grids with differing *x*, *y*, and *z* resolutions and multiple time steps. The user interactively selects the dependent variable to be contoured, the dependent variable to be displayed using color on the isosurface, and the threshold value of the isosurface. We can optionally render the surface semitransparently and allow multiple contours to be displayed at once (see Figure 9).

Finite difference grids traditionally have been developed piece by piece using mesh generators in a very tedious and error-prone process. The accuracy of the mesh is of critical importance. We have developed Image, an interactive mesh analysis tool for verifying finite difference grids. These meshes drive electromagnetic simulations that take hours of supercomputer time. The user manipulates the meshes interactively, choosing slice planes and various components of the mesh for display, and querying for the location of a cell in the three-dimensional mesh using a mouse.

We developed a technique similar to the one described in Upson and Keeler[8] to render semitransparent volumes of three-dimensional data using our volume visualization tool, Cell-Tracer. Ray tracing is used to approximate the rendering integral equation[9] that must be solved for each cell in the grid to "see through" a volume. This powerful technique allows the engineer to browse through the interior of the three-dimensional object and to view the surface of the object.

The AVRP software system is supported on a wide variety of vendor hardware used by the Engineering Department at LLNL. Currently, AVRP tools are available on Silicon Graphics, DEC, Sun, and Stellar platforms. Our philosophy is to port to the low-level graphics library on each system, thereby providing the greatest performance at the lowest software cost.

The scientific visualization challenge is to develop a set of visualization tools that are interactive, provide qualitative and quantitative information, are easy to use, offer new ways of visualizing data, allow the scientist to be more productive, and are available on any workstation. We have described two different attempts at providing these tools. The GWS philosophy is to base everything on standards and require high-performance equipment. The AVRP philosophy is to research and develop new visualization techniques and deliver them to engineers and scientists.

In the future computing environment of LLNL, a scientist will expect to access graphics capabilities in three different modes: graphics created on supercomputers and displayed in the office, graphics created on workstations, and graphics distributed between supercomputers and workstations. To provide these capabilities to each scientist, advances in network technology and workstation performance are essential to maintain the required high throughput rates. In the research areas, we intend to integrate all the AVRP tools so that we have a common environment through a language-based interpreter. We plan to research massively parallel visualization algorithms based on unique and specialized LLNL-developed hardware.

All of our efforts could be enhanced by

scientists and vendors standardizing on a multivariate data representation, to include geometry, topology, and other data objects as NCSA has suggested.[10] Visualization tools built on this standard could be integrated into any environment that makes use of the standards. Visualization tools can aid scientists in gaining insight and productivity and, thus, improve the quality of their science.☐

Acknowledgments

The Magic project was funded and supported by the Defense Sciences Department, A Division, and the Livermore Computer Center at LLNL. The Magic project team consisted of Joe Requa, Jeff Rowe, Bob Lockhart, Dale Nielsen, Pete Keller, Lance Sloan, Dave Wiltzius, Keith Johnson, Dave Handeli, and Pat Gray. Neale Smith provided MacMagic support. The GWS project was funded and supported by the Defense Sciences Department, B Division, and the Livermore Computer Center. The GWS project team consisted of Jeff Long, Gene Cronshagen, Nancy Johnston, Carol Hunter, Linnea Cook, Mike Anastasio, and Bruce Goodwin. The Engineering Department AVRP team consisted of Brian Cabral, Mike Allison, Chuck Grant, Bob Schectman, Mark Wagner, and Bryan Lawver. Jeff Rowe created Figure 1; Pete Keller created Figure 2; Jeff Long created Figures 3, 4, 5, and 6; Brian Cabral created Figures 7, 8, 9, and 10. We appreciate the efforts of Derek Hendry, Gary Laguna, and Rebecca Springmeyer in reviewing this article.

References

1. B. McCormick, T. DeFanti, and M. Brown, "Visualization in Scientific Computing," *Proc. ACM SIGGraph,* Vol. 21, No. 6, Nov., 1987, p. 3.
2. C. Hunter et al., "Requirements for Octopus Display Replacement," Lawrence Livermore National Laboratory, Livermore, Calif., UCRL-53886, Sept. 12, 1988.
3. J. Rowe, "Magic," *Tentacle,* Vol. VII, No. 5, Computation Dept., Lawrence Livermore National Laboratory, May 1987, pp. 1-7.
4. J. Hardin, A. Stowell, and F. Bond, "Creating Tools for Science," *NCSA access,* Vol. 3, No. 1, Jan.-Feb., 1989, pp. 8-10.
5. C. Hunter, "New Graphics for Cray Users," *Proc. Cray User Group,* Spring, 1987, pp. 168-171.
6. L. Cook, "Visualization, Prototyping, and Scientific Togetherness," *Unix Review,* Mar., 1989, pp. 48-50.
7. W. Lorenson and H. Cline, "Marching Cubes: A High-Resolution 3D Surface Con-

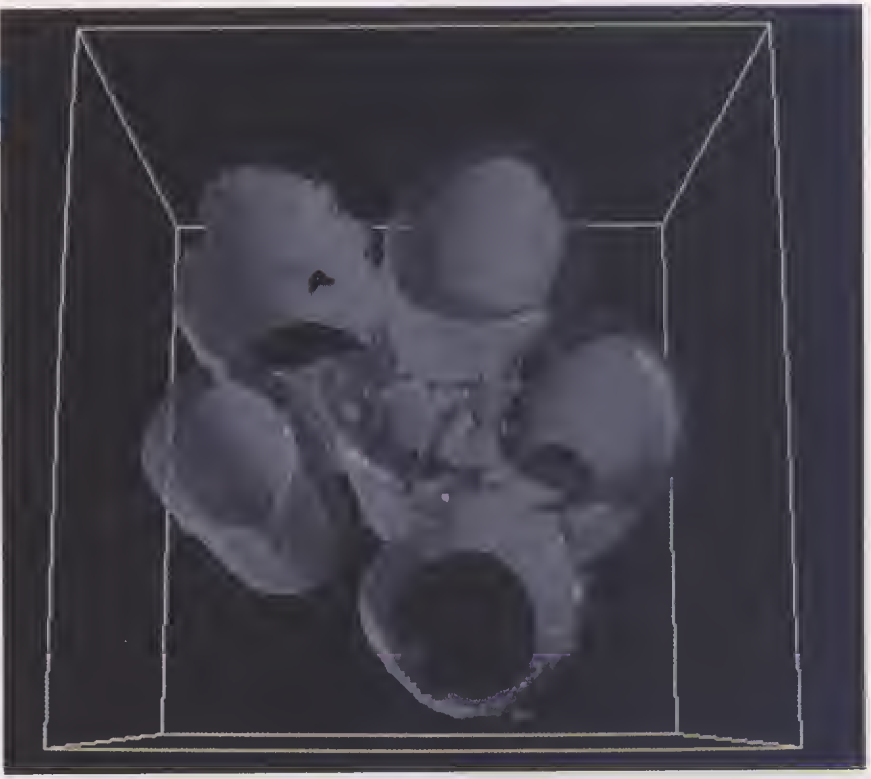

Figure 9. Isosurface rendering of a three-dimensional reconstructed image of data from a nondestructive evaluation using a neutron transport.

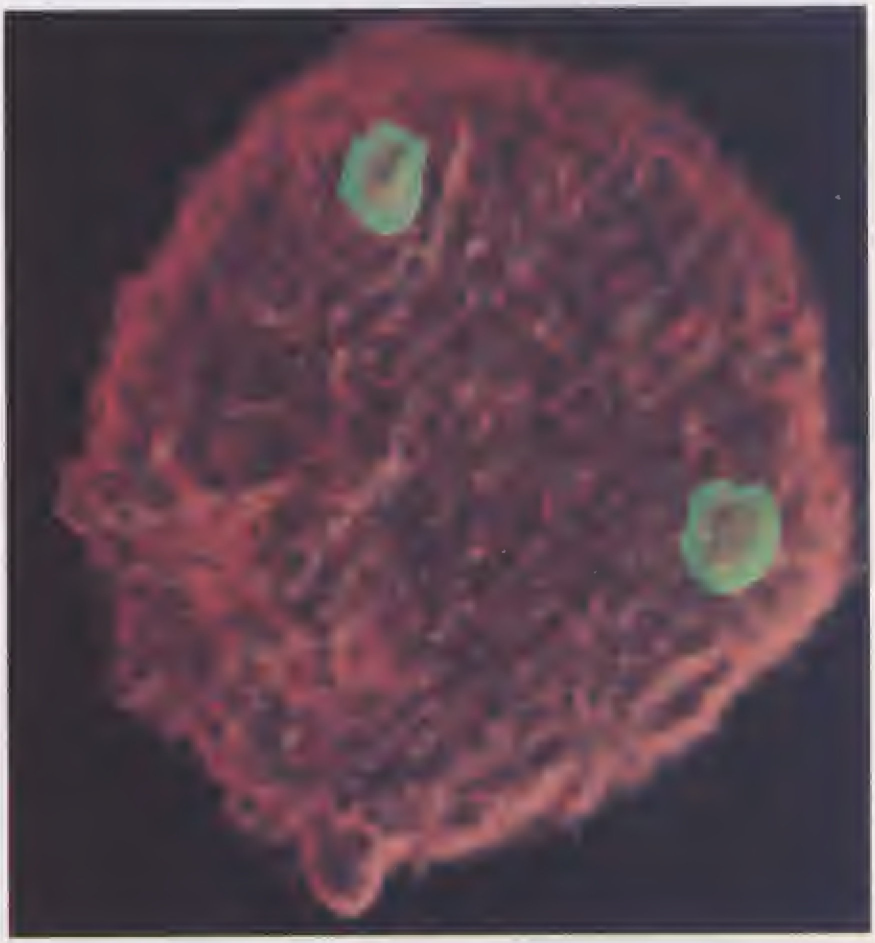

Figure 10. Semitransparent volumetric rendering of the nucleus of a human cell. The red clumps are chromosome groupings within the cell nucleus.

struction Algorithm," *Proc. SIGGraph 87,* pp. 163-169.
8. C. Upson and M. Keeler, "V-Buffer: Visible Volume Rendering," *Proc. SIGGraph 88,* pp. 59-64.
9. J. Kajiya, "The Rendering Equation," *Proc. SIGGraph 86,* 1986, pp. 143-150.
10. M. Folk, "HDF: A File Format for Sharing," data link, Nat'l Center for Supercomputer Applications, Univ. of Illinois at Urbana-Champaign, Jan.-Feb., 1989, pp. 1-9.

Brian Cabral is a computer scientist in the Engineering Department at Lawrence Livermore National Laboratory and is project leader of the Advanced Visualization Research Project. His interests include computer graphics algorithms and systems, computational geometry, operating systems, computer architectures, and computer languages. Cabral received a BS from California State University, Stanislaus, in 1983 and an MS from the University of California at Davis in 1988, both in computer science.

Carol L. Hunter is the graphics coordinator and the mainframe graphics group leader for the User Systems Division of the Livermore Computer Center at Lawrence Livermore National Laboratory. She is responsible for the long-term planning, implementation, and integration of graphics hardware and software to support the 2,000 supercomputer and workstation users of the center. Previously, she was the graphics workstation project leader, a team leader, and a computer scientist for the Nuclear Software Systems Division at Livermore. She is a member of SIGGraph and ACM.

Cabral can be contacted at Lawrence Livermore National Laboratory, Mail Stop L-156, PO Box 808, Livermore, CA 94550, or electronically via cabral@hummer.llnl.gov.

Scientific Visualization and the Rivers Project at the National Center for Supercomputing Applications

Robert B. Haber, National Center for Supercomputing Applications

The rapid growth of large-scale computing in the basic sciences and the steady accumulation of high-bandwidth data sources (radio telescopes, medical scanners, etc.) are largely responsible for the recent interest in scientific visualization as a computational technology. This article reviews activities in scientific visualization at the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign.

NCSA is one of five National Science Foundation-sponsored supercomputing centers that serve the national research community in computational science and engineering. The center began operations in 1985 and now serves a national user base of more than 2,500 academic researchers in virtually all areas of science and engineering. NCSA has also established an industrial partners program, which involves five major US corporations. These companies house individual researchers at the center and have network access to the supercomputers at their home offices. From its inception, NCSA has promoted the development and use of scientific visualization as a necessary ingredient for computational science and engineering.

NCSA's primary computing resources are a Cray X-MP/48 with a 128-megaword solid-state storage device and a Cray-2S/4-128. The larger systems are supplemented by an Alliant VFX-80, an Alliant FX-8, and a large number of workstations and personal computers. In addition to delivering reliable service to its users, NCSA includes in its mission the education and training of new supercomputer users, expansion of the role of computational science and engineering in academic and industrial research, and development of advanced computational technologies to support research.

Overview of NCSA activities in scientific visualization

NCSA's user community is multidisciplinary and contains a wide spectrum of expertise in computational technology and visualization techniques. As a result, NCSA has developed a variety of services and initiatives to provide visualization capabilities to its users. The following sections contain a brief review of visualization-related activities at the center, and the research and development activities of the Rivers Project.

Visualization and Media Services Group. The Visualization and Media Services Group contains two recently merged parts: Scientific Media Services and the Scientific Visualization Project. SMS produces film and video output of graphics generated during scientific research by NCSA researchers and computer users. Automated procedures have been implemented to produce hard copy from user-generated image files stored in standard formats. These files might be created directly by users (perhaps with guidance from consultants) with available graphics applications packages—such as NCAR (National Center for Atmospheric Research) graphics, Movie.BYU, and NCSA Image—or by NCSA's visualization staff.

The images are first displayed on a high-resolution Raster Technologies frame buffer attached to the Alliant VFX-80, converted to standard television signal NTSC (National Television System Committee) format, and then transferred to videotape. The problem of producing smooth video animation from images, which might take minutes per frame to generate, is solved using an Abekas digital video recorder, a device commonly used in entertainment, special effects, and commercial animation applications. Recently, real-time scan converters were added to transform high-resolution RGB (red, green, blue) video signals from the frame buffers to NTSC format. Another recent equipment upgrade includes an Abekas A60, supporting direct-image file transfer from Ethernet and digital component video recording (this eliminates the need to convert to the lower-quality composite NTSC video format).

The Scientific Visualization Project is staffed by computer graphics and visualization experts who produce custom, high-quality visualization animations for scientists selected through a peer review process. This service is not automatically available to all NCSA users. SVP typically generates images on the Alliant VFX-80 (or a graphics workstation) using a rendering package from Wavefront Technologies (or custom software developed at NCSA).

In a typical production cycle, the scientist confers with the SVP staff to discuss alternative methods to visualize the simulation data. SVP then develops an animation storyboard to plan the appearance, motion control, and sequence of the animation. The thousands of frames needed to create a typical animation are produced under the control of an automated batch rendering control system. SMS then transfers the images to videotape or slides. A typical image frame produced by this method appears in Figure 1.

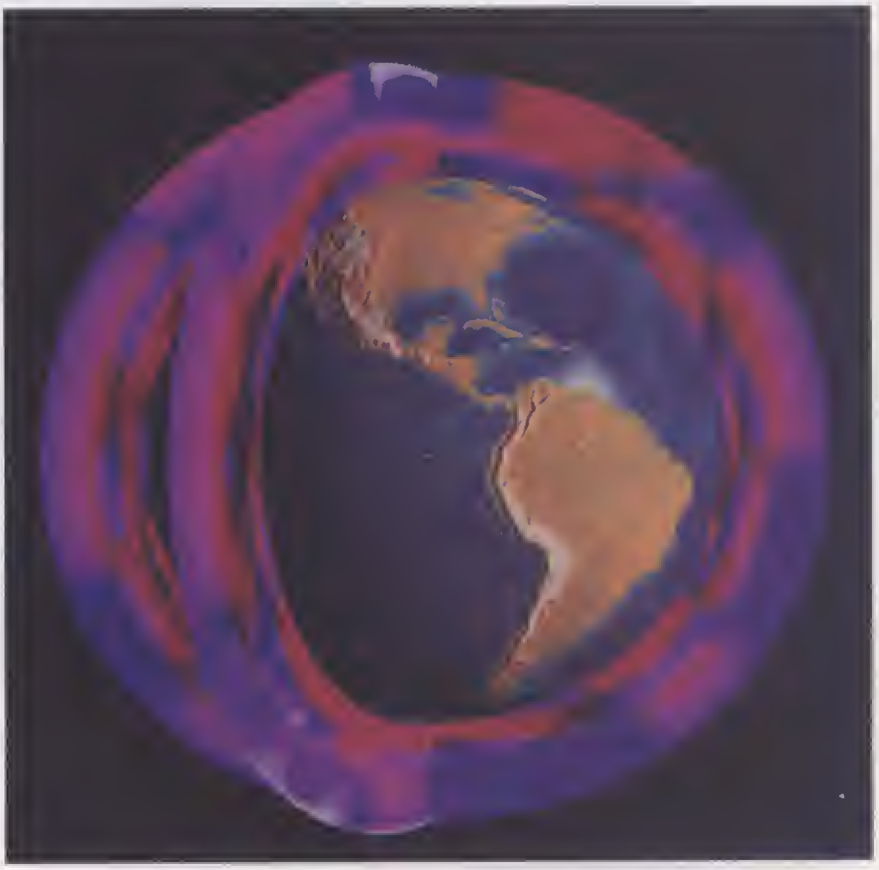

Figure 1. Visualization of global climatic effects of increased greenhouse gases. (W.M. Washington and T.W. Bettge of NCAR performed the simulation; J.B. Yost of NCSA performed the visualization.)

Software Development Group. The Software Development Group (SDG) generates original software packages for computational science. In many cases, prototype software systems, developed in other parts of NCSA, are converted into useful products with reliable software engineering, full documentation, and enhanced features. The software is developed to run on equipment from several vendors (such as Apple Computer, IBM, and Sun Microsystems).

To date, much of this effort has been targeted at low-cost hardware platforms to provide improved computing environments for NCSA's many users. SDG software may be obtained free of charge via anonymous FTP (file-transfer protocol) across the network, or for the cost of distribution media and handling via catalog order.

SDG software is a key element in NCSA's efforts to deliver visualization capabilities to remote users. NCSA Telnet, the first SDG product to emerge (with more than 20,000 users in the US, at last count), provides multiple simultaneous session connection over TCP/IP networks, a built-in FTP server for file transport, and VT102 and Tektronix 4014 graphics emulation on IBM and Apple Macintosh personal computers.

There is also a growing suite of visualization software tools designed to work together to provide a comprehensive solution. NCSA Image and NCSA ImageTool provide color raster image display and manipulation capabilities to users with Macintosh and Sun workstations (see Figure 2). NCSA DataScope enables the scientist to simultaneously display floating-point arrays and images generated from those arrays. NCSA PalEdit supports interactive construction, editing, and manipulation of color lookup tables for raster displays, and is designed to work with the other color display programs. NCSA CompositeTool allows the scientist to composite, overlay, annotate, and prepare presentation quality images for slides and prints. Current work in SDG is aimed at developing X Window System support for the entire software suite to significantly expand the base of supported hardware. Work is also beginning on software tools for 3D and volumetric rendering.

User training, consulting, and education. NCSA places a high priority on training and education in its mission to expand the use of supercomputing and visualization. NCSA hosts frequent workshops and seminars throughout the year, some targeted specifically at scientific visualization. Intensive training is provided to consultants who serve as local resources at the 60 colleges and universities participating in the center's Academic Affiliates Program. A visitors' program supports individual researchers who wish to visit NCSA for direct interaction with consultants and the visualization staff.

Recently, NCSA established the Renaissance Experimental Laboratory, a classroom for supercomputing and visualization education equipped with 20 personal IRIS workstations donated by Silicon Graphics. Classes in visualization in the REL are offered through the University of Illinois Computer Science Department. Other courses emphasize applications of visualization in mathematics, computational chemistry, biophysics, art, and design.

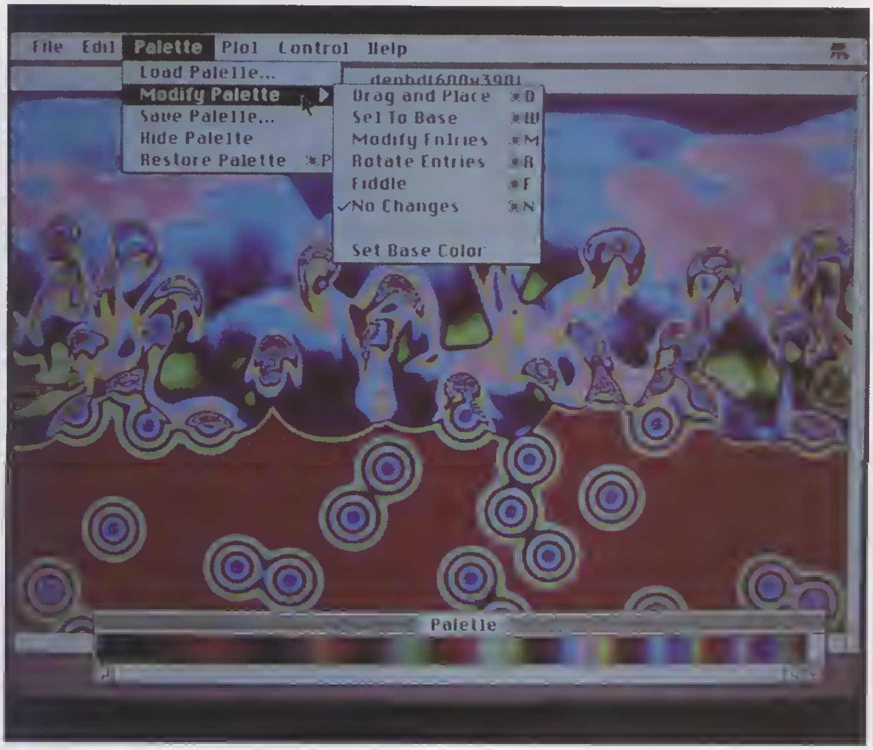

Figure 2. **Visualization using NCSA Image of an astrophysics study of supernova-induced shock waves accelerating clouds in the interstellar medium. (M.L. Norman, J.R. Dickel, M. Livio, and Y.-H. Chu of NCSA and the Department of Astronomy of the University of Illinois.)**

The NCSA Rivers Project

During the fall of 1987, NCSA initiated a research project to explore several emerging computational technologies that promise to have a significant impact on the scientific computing environment. This program became known as the Rivers (Research on Interactive Visual Environments) Project. Specifically, Rivers is developing hardware and software systems for extending high-end 3D visualization from a batch process to an interactive process, and for visualization-based interactive steering of supercomputing simulations in a high-performance distributed environment.

The process of visualization. A structured view of the process of visualization motivates the Rivers hardware and software architectures. Visualization can be viewed as a series of transformations that convert raw simulation data into a displayable image. The goal of these transformations is to convert the information to a format amenable to understanding by the human perceptual system while maintaining the integrity of the information. Three major transformations occur in most visualization procedures:

- The *data enrichment and enhancement* step operates on the raw data provided by the simulation and modifies it in one or more ways to derive data for subsequent visualization operations. Examples include data filtering to remove noise (numerical or observational), calculation of gradients in field solutions, and interpolation to a regular grid.
- *Visualization mapping* can be viewed as an operation that constructs an imaginary object, here called an abstract visualization object or AVO, from the derived data produced by the enhancement and enrichment operations. Typically, this involves a series of mappings of the simulation data into the attribute fields that describe the AVO (geometry, time, color, transparency, luminosity, reflectance, surface texture, etc.). A common example is the mapping of scalar field data into height or color.
- *Rendering* is the last transformation in the visualization process. The rendering mapping operates on the AVO to produce a displayable image. Rendering can involve familiar operations from computer graphics and image processing, but might also involve new algorithms that respond to new requirements of scientific and engineering visualization, such as volumetric rendering.

Important conclusions can be drawn from this model of visualization:

(1) The process can be conveniently segmented into modular operations.
(2) In many cases the modules can be made generic and shared across several science disciplines.
(3) A specific visualization process can be defined as a dataflow graph with mapping modules as nodes (the graph can

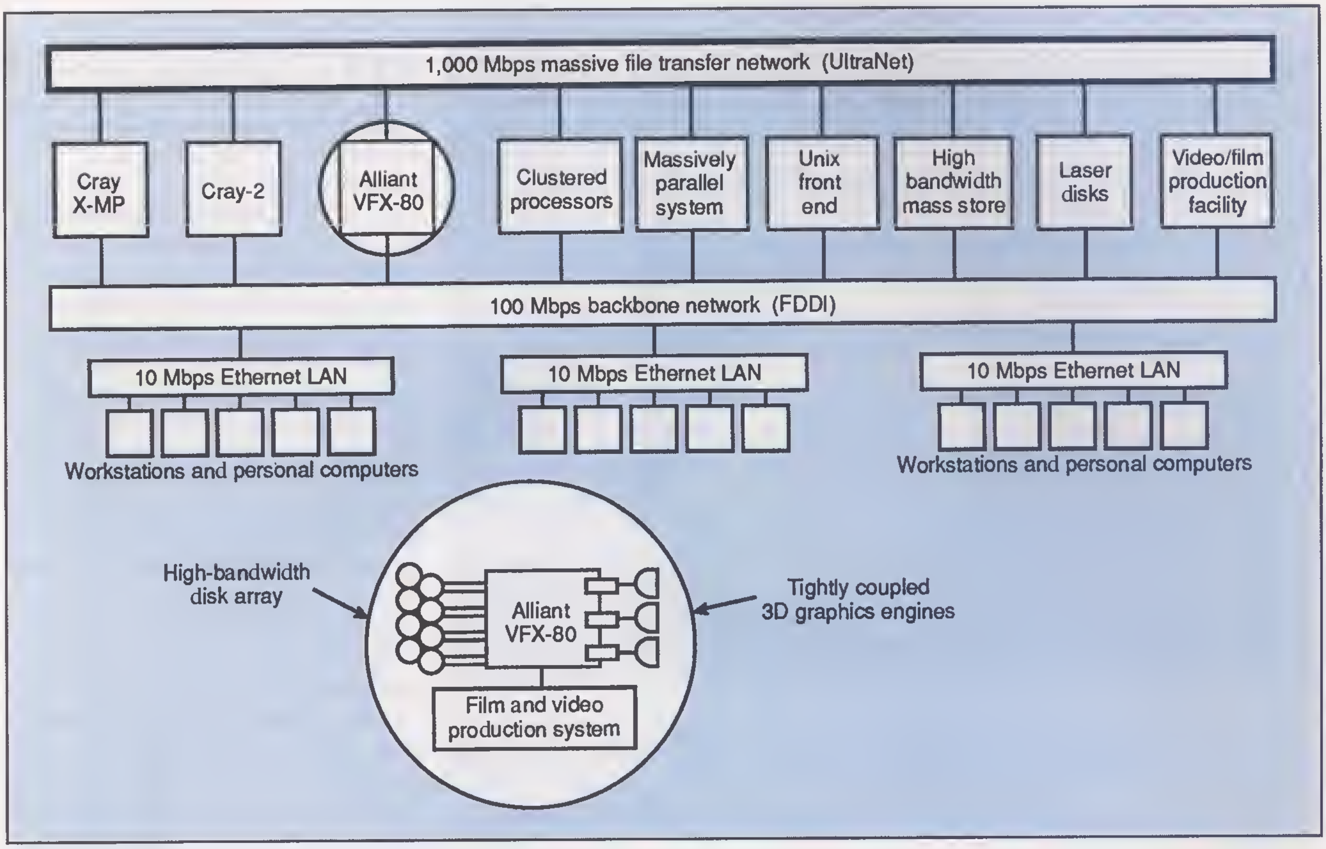
1,000 Mbps massive file transfer network (UltraNet)
Cray X-MP
Cray-2
Alliant VFX-80
Clustered processors
Massively parallel system
Unix front end
High bandwidth mass store
Laser disks
Video/film production facility
100 Mbps backbone network (FDDI)
10 Mbps Ethernet LAN
10 Mbps Ethernet LAN
10 Mbps Ethernet LAN
Workstations and personal computers
Workstations and personal computers
High-bandwidth disk array
Alliant VFX-80
Tightly coupled 3D graphics engines
Film and video production system

Figure 3. Hardware architecture for distributed simulation and visualization specifies the NCSA three-tier computing strategy.

be reconfigured to define alternative visualization methods, and the system can be extended by adding new modules).

(4) Interactive control over mapping parameters is desirable to fine-tune a given idiom.

(5) Visualization has a pipeline structure that can be exploited in hardware and software design to obtain improved performance.

The Rivers software design. A modular software system is under development with the above model in mind. Key principles of the Rivers software design are

- software modules should support distributed execution across multiple, heterogeneous hosts;
- appropriate user interfaces should be available for defining configurations and for adjusting parameters;
- low-level issues such as graphics languages and communications protocols should be hidden from the user (a layered design);
- software standards (X.11, PHIGS+, Unix, etc.) should be used to the greatest extent possible; and
- software components should be designed for compatibility with high-performance systems for interactivity.

The Rivers hardware architecture. The pipelined structure of the visualization process requires careful attention to balance in hardware system design. In addition to raw graphics performance, a balanced visualization system must provide commensurate sustained bandwidth between distributed hosts (that is, the network); between mass storage systems and host memory; and between host memory and special-purpose graphics accelerators. Multiprocessor, shared memory architectures are also attractive since they map well onto the pipelined structure of visualization. Since visualization includes a considerable amount of computation-intensive calculation in the early stages of the pipeline, powerful general-purpose processors are needed in addition to powerful rendering engines. At present, the general-purpose computation is often the main bottleneck.

Figure 3 diagrams the distributed hardware system for visualization now under development within Rivers and other segments of NCSA. A key theme is the use of dual, three-tier schemes for both the computing and network strategies. Middle-tier systems, such as the Alliant VFX-80, provide higher performance capabilities that can be shared by a group of users without the compromises necessary with national resources like the true supercomputers. For example, the VFX-80—with a closely

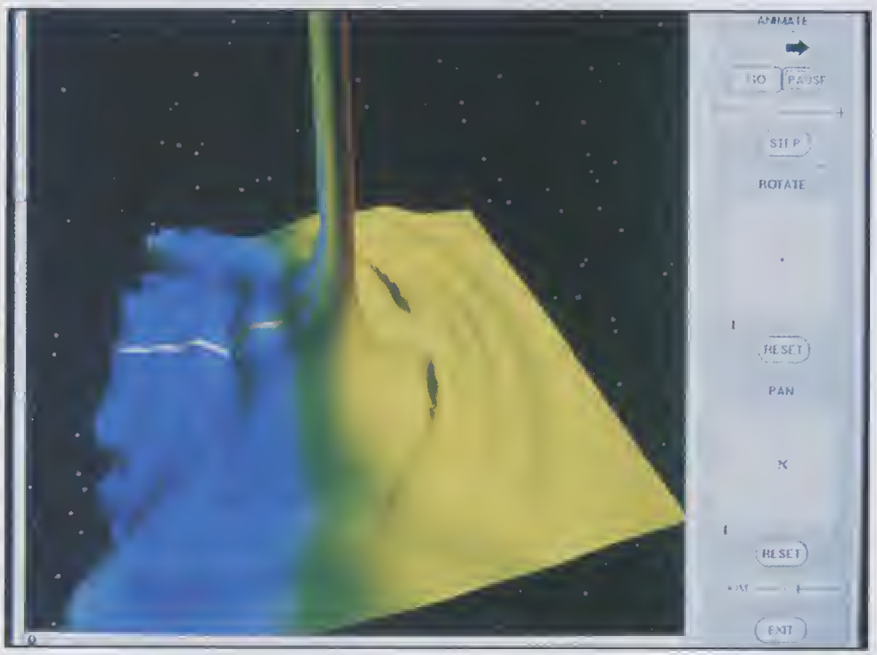

Figure 4. Visualization of rapid crack growth in a brittle aluminum plate using the Rivers polygon animation tool. (H.M. Koh, H.S. Lee, and R.B. Haber, NCSA and Departments of Theoretical and Applied Mechanics and Civil Engineering, University of Illinois (simulation); S. Chall, R. Idaszak, and P. Baker, NCSA (visualization).)

coupled visualization system, a connection to the Ultra network, and a disk array mass storage system (planned)—provides a shareable, high-performance visualization capability that is too expensive to reproduce on every desk top.

All high-level machines, including the supercomputers and middle-tier processors, are connected by the highest speed network available (NCSA installed an UltraNet system in December 1988, with a peak bandwidth of 100 megabytes per second). This network will be reserved and optimized for massive data transfer using large packet sizes (for real-time visualization, distributed processing, and accessing the shared high-performance mass storage system).

A middle-tier network forms the backbone for handling messages, smaller data packets, and conventional network traffic between the large machines and the bottom tier networks connecting the workstations. At present, a 50-megabit-per-second HyperChannel network and an 80-megabit-per-second ProNet-80 network fulfill this function. (The FDDI, or Fiber Digital Data Interface, standard is expected to take on the middle-tier function when it becomes available). The objective here is to segregate the massive data transfers from terminal keystrokes and smaller messages to provide the best efficiency and performance.

Hardware engines for accelerating the rendering operations are a key part of the plan. At present, the Rivers Project includes a variety of graphics accelerators for polygon rendering (Raster Technologies GX4000 boards on Sun-4 platforms, the Alliant VFX-80, SGI workstations) and for volumetric rendering (the Sun TAAC-1 application accelerator).

Figure 4 shows a prototype three-dimensional, interactive animation tool that the Rivers group developed. The tool is capable of animating complex polygon data sets in real time on the workstation or VFX-80 screen. Although relatively simple, the tool has been used to view data from a variety of disciplines—solid mechanics, atmospheric sciences, and astrophysics.

The visualization approaches described above vary considerably in cost and applicability to the requirements of various computer users. The SDG software suite runs on inexpensive platforms and exemplifies visualization technology immediately practical at almost any computer site. The SMS production facility requires staffing and capital expenditures that imply a reasonably large user base. Lower cost visualization facilities can be assembled by trading off performance, quality, and flexibility. Also, several computer vendors are working on new products that provide integrated visualization systems at lower price levels. Thus, this technology should become accessible to a larger user group with time.

High-quality visualization production by graphics experts using flexible software rendering systems continues to produce impressive results. The expense of this approach remains high, due to the labor-intensive nature of the process and the relative scarcity of visualization experts. High-performance, interactive visualization as in the Rivers project is also currently expensive due to equipment costs. While the necessary graphics performance is now available in some high-end graphics workstations, it will take several years for the cost of balanced mass storage and network systems to become affordable for single users. Until then, the techniques used in Rivers will be suitable for shared, department-level facilities. Therefore, the Rivers visualization system can be viewed as a development prototype for affordable, workstation-based systems that might appear in three to five years.

Recent experience in the Rivers project has demonstrated that very-high-performance systems can be built using hardware and software standards. Standards are particularly important in helping to amortize the significant software development effort needed to make high-level visualization tools available to applications users.

The emerging HSC (high-speed channel) standard interface will facilitate the integration of the high-performance, distributed systems needed for real-time interactive visualization. National gigabaud networks (proposed for the mid to late 1990s) will allow remote access to interactive supercomputing and visualization. ☐

Acknowledgments

Many talented individuals have contributed to the success of the visualization program at NCSA. It is impossible to identify all of the individual contributors here. However, I want to credit Matt Arrott, Polly Baker, Dan Brady, Steve Chall, Donna Cox, Stefen Fangmeier, Joe Hardin, Ray Idaszak, Tim Krauskopf, Dave McNabb, Nancy St. John, and Craig Upson for their leadership and contributions to visualization development at NCSA.

Various aspects of the visualization work at NCSA are supported by the National Science Foundation, the state of Illinois, Alliant Computer Systems, Apple Computer, Cray Research, Silicon Graphics, Sun Microsystems, Ultra Network Technologies, and NCSA's corporate partners.

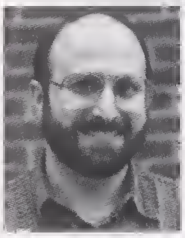

Robert B. Haber is an associate professor of theoretical and applied mechanics and civil engineering at the University of Illinois at Urbana-Champaign. He is also a research scientist at the National Center for Supercomputing Applications, where he directs the Rivers Project. His research interests include finite element methods in computational solid mechanics, computer-assisted design and optimization, and applications of computer graphics visualization to computational science and engineering.

Haber received a bachelor of architecture degree in 1977 and a PhD in civil engineering in 1980, both from Cornell University. He is a member of the American Society of Civil Engineers, the American Academy of Mechanics, the International Association for Computational Mechanics, ACM (SIGGraph), the American Association for the Advancement of Science, and Sigma Xi. He is a Cray Research Affiliate, and president of Allus, the Alliant Users' Society.

Haber can be reached at the National Center for Supercomputing Applications and the Departments of Theoretical and Applied Mechanics and Civil Engineering, University of Illinois at Urbana-Champaign, 254 Computing Applications Building, 605 E. Springfield Ave., Champaign, Illinois, 61820.

Scientific Visualization in Computational Aerodynamics at NASA Ames Research Center

Gordon V. Bancroft, Todd Plessel, Fergus Merritt, and
Pamela P. Walatka, Sterling Software
Val Watson, NASA Ames

At NASA Ames Research Center, scientists are using the Numerical Aerodynamic Simulation (NAS) facility to perform experiments in computational aerodynamics.[1] Scientists model complex fluid mechanics problems using supercomputers and new numerical algorithms. The computer-generated models show vehicles and the air space surrounding them. Physics formulas are used to simulate the flow of air about the vehicle. Scientists use high-performance computer graphics workstations to view and, in some cases, animate these simulations.

The high-performance graphics workstation is connected to a supercomputer with a high-speed channel. Postprocessing (processing performed after the simulation is complete) is by far the most commonly used method for visualization of computational aerodynamics. A flow simulation is executed on a supercomputer, and then the results of the simulation are processed for viewing.

Tracking and steering involve interaction. Tracking allows the scientist to track his or her flow solution so it can be stopped and, perhaps, restarted should undesirable results be obtained. Steering goes one step further and allows the scientist to change parameters of the flow solver as it executes. Examples of this might be increasing the angle of attack of an aerodynamic body, stretching a grid to move points, or perhaps increasing the Mach number. These are more desirable methods than postprocessing, yet much less common given the current state of supercomputer and workstation evolution and performance. Distributed processing is the term used when the process of visualizing the solutions is distributed between the workstation and the supercomputer.

The software cycle for fluid dynamics simulation begins with the creation of a grid that describes the walls of a vehicle and the surrounding space. Flow-solver programs (that is, algorithms derived from physics formulas for the flow of air) are then executed (typically on supercomputers) to generate flow-field solutions. The raw data from these solutions consists of density, momentum vector, and total energy per unit volume specified at each grid point in the computational domain. A typical computational domain may contain one million grid points. This raw data must be converted to a scene depicting the physics in a manner the scientist can easily interpret. Color and visual cues (shading, animation, etc.) are used to demonstrate the physics of the particular result.

Visualization requirements

The views of the simulation portrayed by the computer graphic workstations must be 3D because visualization of the interrelated flows of all three dimensions simultaneously is important. The scenes must be dynamic for the time-variant features of the flow fields to be understood. Although the motion need not be real time, the motions must be rapid enough to give a proper understanding of the dynamic features of the flow.

The flow fields typically have a large range of scales. Therefore, the scientist must be able to zoom into a region of small-scale features and zoom back out to view the overall flow field. Furthermore, the scenes should be high definition to contain adequate detail at all scales.

The scenes should simultaneously contain solid body objects, such as an aircraft (with hidden surfaces removed), and points or lines (such as lines representing the paths of tracer particles inserted into the flow field). As the scenes evolve in time, illustrating the flow dynamics (for example, the movement of tracer particles), the viewing position must be simultaneously changeable in real time (as the flow evolves) to maintain the best view or to get a different perspective.

Dynamic change of the viewing position is one of the best cues for enhancing the 3D aspects of the scene. In addition, new visualization effects, such as ribbon traces, smoke, shading of function-mapped parts, antialiasing, variable transparency, volume visualization, and stereo, are being requested by scientists to further enhance their ability to analyze flow fields.

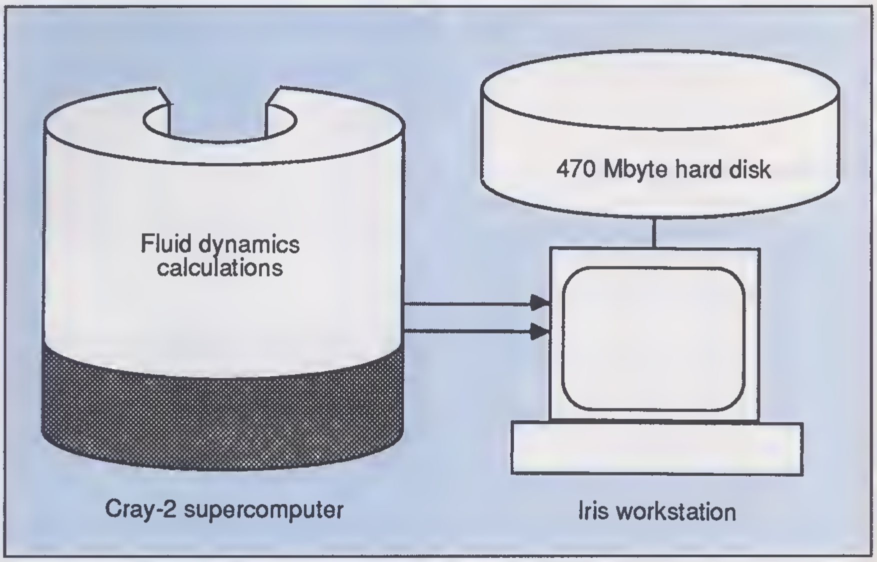

Figure 1. Scientists use the supercomputer to generate solutions to aerodynamic problems, with five variables (density, x-, y-, and z-momentum, and energy) for each of a million or more grid points. The HyperChannel link is used to move the data to the workstation, where the numbers are turned into graphics. Many functions, such as pressure, shock wave, or particle trace, can be plotted from the solutions. The interactive visualization greatly enhances the understanding of the solutions.

Hardware

The hardware currently in use for computational fluid dynamics at NASA Ames includes supercomputers, workstations, digital video recording equipment, film recorders, still cameras, color printers, digital scanners, and color copiers.

The calculations to generate the flowfield solutions are done on the supercomputer (see Figure 1). The visualizations—animations and scenes—can be created in three ways:

(1) Transfer the whole solution file (containing the solution at each grid point) to the large disk on the workstation and generate the scene on the workstation.
(2) Produce graphics files on the supercomputer and transfer these graphics files to the workstation.
(3) Create the scene using the supercomputer interactively while viewing the scene on the workstation.

The key features of a workstation are its rapid 3D transformation speed (for changing the viewing position), its high-definition display, and its rapid display creation speed. On the current workstation, a Silicon Graphics 3130, 3D coordinates can be transformed at a rate of 80,000 coordinates per second. The display has high spatial resolution (1,024 pixels horizontally by 768 pixels vertically) and high color resolution (24 color planes giving more than 16 million simultaneous colors). Scenes with a very simple solid object and thousands of lines or points can be generated at a rate of more than 10 per second—a rate that provides satisfactory motion for understanding dynamics.

Table 1. Features of current and next-generation workstations.

Feature	Iris	Future W/S
CPU	2,600 Dhrystones 0.1 Mflops	20,000 Dhrystones 2.0 Mflops
Memory	4 MB	16 MB
Bus	Multi (2 MB/sec)	VME (>20 MB/sec)
Disk	474 MB	1,000 MB
3D trans.	80 K/sec	800 K/sec
Filled polys	16 K/sec	200 K/sec
Z-buffered	0.5 K/sec	100 K/sec
Dsply mem.	32 × 1 K × 1 K	76 × 1 K × 1 K

(Polygons 400 pixels; z-buff polygons are shaded)

The workstation contains a z-buffer[2] for hardware implementation of hidden-surface removal. However, many seconds are required to create scenes of typical aerodynamic vehicles if the z-buffer and Gouraud shading[5] are used. Therefore, these scenes must be recorded on videotape or 16-mm movie film to allow satisfactory viewing of the dynamics.

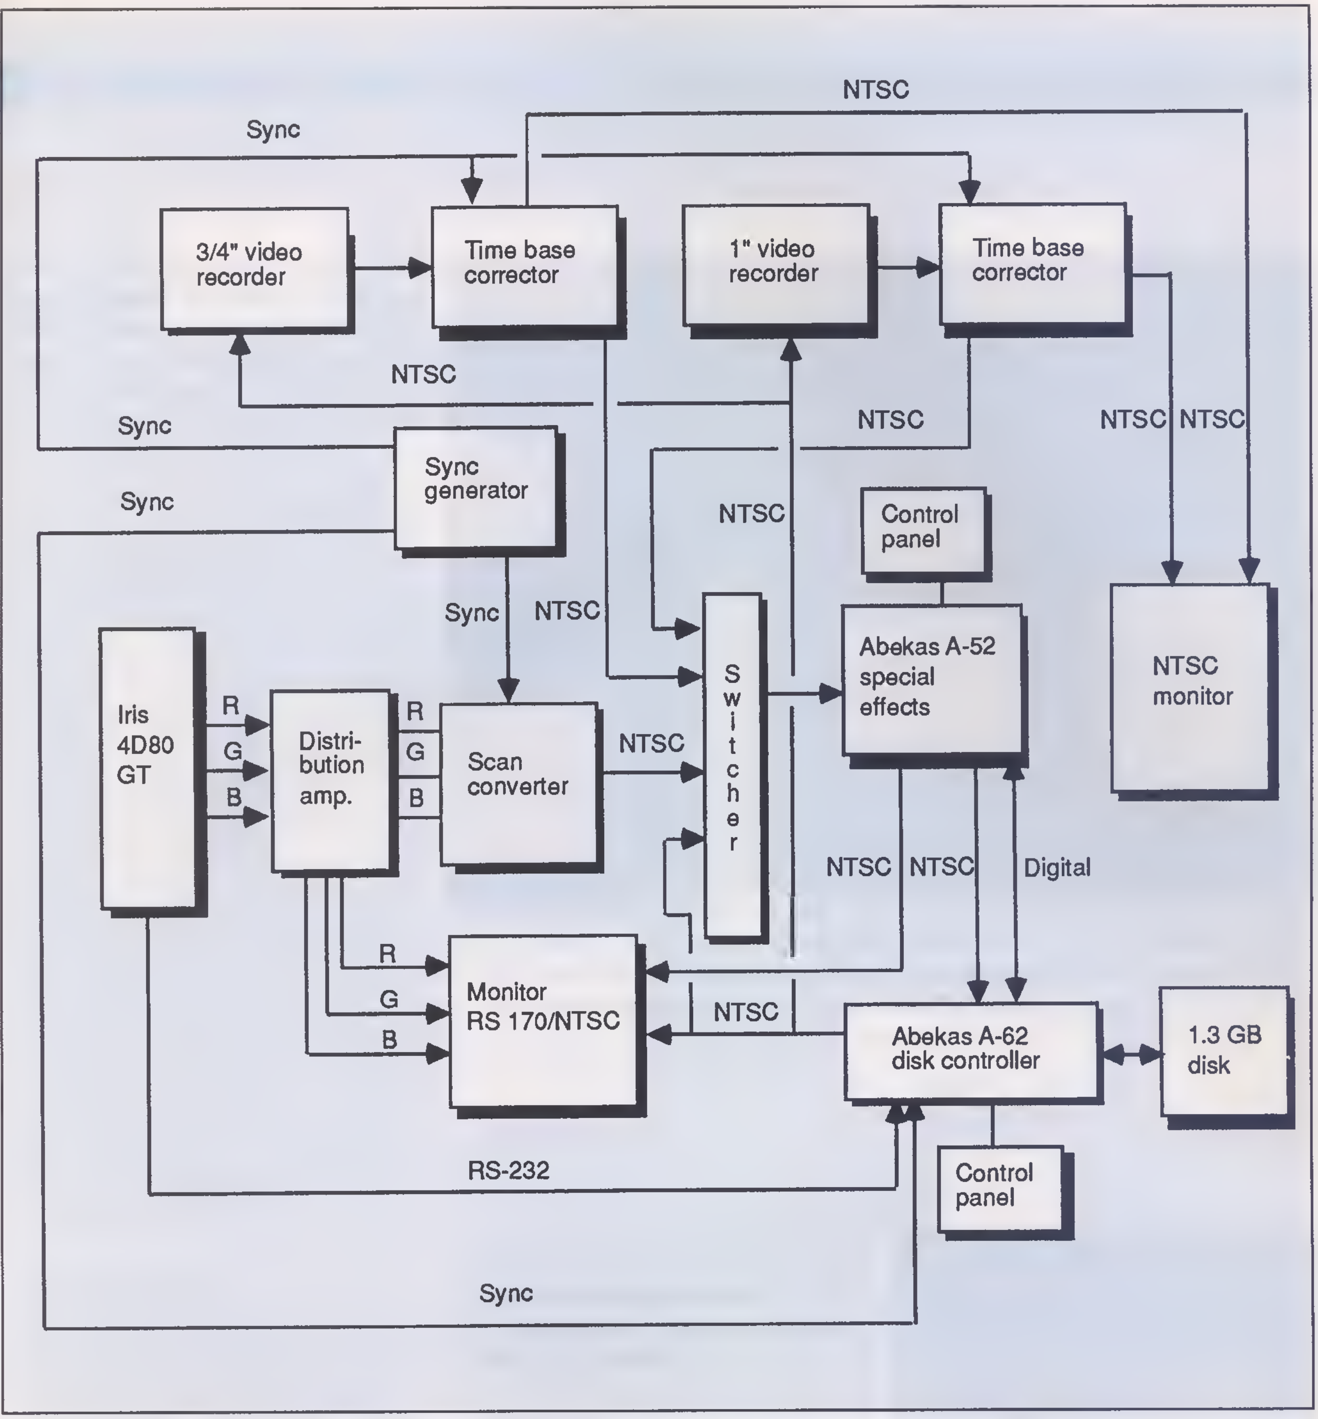

Figure 2. This Abekas video recording system diagram is used to animate the aerodynamic scenes. Animation enhances understanding of the flow dynamics by showing the scenes from different viewpoints or for different times. Simple animations can be done on the workstations, but group presentations or complex scenes require video or film recording.

The next generation of workstations, to be available to NASA scientists before the end of 1989,[3] is expected to meet the requirements for dynamic viewing, as indicated in Table 1. Benchmark software has been developed at NASA Ames to test what kind of graphics performance can be expected from these next-generation workstations (this software is available at no charge through the authors from NASA Ames as part of our graphics development effort).

Currently, animations of the scenes are recorded to videodisc or videotape with the hardware shown in Figure 2. The high-definition scene is digitally sampled by a scan converter to a lower resolution RS170a format. This format can be encoded from separate red, green, and blue components into the standard single composite NTSC (National Television Standards Committee) signal used by standard video recorders and players. As each frame is displayed, control information tells the Abekas to record. It then stores the frame as digital NTSC.

This process occurs at real-time rates, so

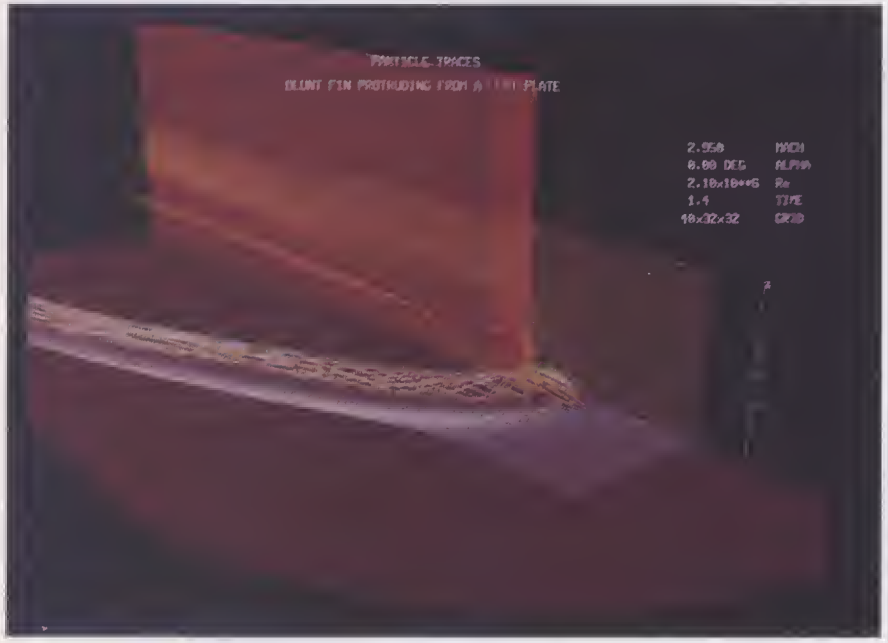

Figure 3. This is an example of a scene created with Plot3D, one of the postprocessing software packages used for visualizing computed aerodynamic solutions. The object is a blunt fin (simplified abstraction of a wing protruding from a wall); the rake-like lines are particle traces, indicating the streamlines of the flow of air.

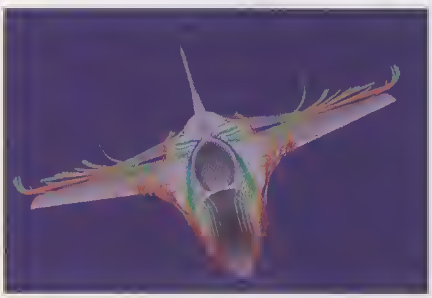

Figure 4. GAS (Graphics Animation System) was used to assemble the parts in this computer-generated scene. The particle traces were computed with Plot3D, the smooth-shaded F16 body was created with Surf. Then, GAS was used to combine the parts and interactively change the colors and viewing angle. GAS is also used to animate scenes.

the time required to record a computer workstation frame is limited by the time it takes to render it. The workstation then continues on with the next frame and repeats the process until the animation is complete. The Abekas uses Winchester disk technology (1.3 gigabytes per 100 seconds of video), allowing stored video to be edited (using the A52 special effects) or the disks to be rerecorded. There are no generation losses within the system due to the digital formatting.

The hardware for recording the scenes on 16-mm film is shown in Figure 3. The Dunn camera is controlled from the workstation using an RS-232 hardware connection and the Graphics Animation System (GAS) software described later in this article. Other Dunn cameras and color printers are used for making stills.

Software

The following are examples of postprocessing software packages in use at the NASA Ames Research Center:

Plot3D accepts as input the flow-field solutions from the supercomputer and creates as output a variety of scenes that can be viewed dynamically with the workstations [4] The software makes extensive use of color and 3D cues (such as shading and perspective as on p. 269 of Foley and Van Dam[2]). A very popular scene is path lines of particles released at selected points inside the flow field (Figure 4). Another example is color mapping on a vehicle surface representing the magnitude of some scalar property on the surface, such as pressure. Plot3D software is available for properly configured Silicon Graphics, Apollo, and Sun workstations; Cray2/Unicos, Alliant/Unix, and Convex supercomputers; and VAX 11/780 minicomputers. Output can be saved to files for use with the GAS animation program.

Surf (Surface Modeler) allows scientists to input grid and solution files and interactively build a 3D model consisting of wireframe, shaded, and function-mapped parts (Figure 5). These parts can be interactively viewed, edited, and output to files that can then be loaded into GAS and animated. Surf has a mouse-driven interface. Color and specular highlighting can be adjusted interactively. Shaded parts are created based on user-specified light sources (up to 20), a viewpoint, and an ambient light level.

GAS (Graphical Animation System) permits the scientist to interactively and dynamically view and combine the 3D scenes created by Plot3D, Surf, or several other graphical packages while simultaneously changing the viewing position within the 3D space. In addition, it permits the scientist to generate an animation sequence with smooth 3D transitions between a series of specified positions. Titles can be inserted, and the resulting "movie" can be recorded automatically to either video or 16-mm film. This software is device specific and runs only on a Silicon Graphics Iris Workstation.

FAST (Flow Analysis Software Toolkit) is a new proposed standard fluid-dynamics graphics environment. The purpose of FAST is to provide the scientist with a single software environment for handling many graphics needs in a way that is quick, powerful, and easy to use. The FAST environment is being designed for the capabilities of the next-generation workstation (see Table 1). The new capabilities of these machines warrant a new approach to building graphics tools.

RIP (Real-time Interactive Particle-Tracer) is an example of a distributed graphics tool and actually consists of two programs that communicate over a high-speed network.^{5,6} One program computes the flow traces from raw data on a supercomputer, and the other program renders these traces for interactive viewing on a workstation. Particle tracing is then interactive; a scientist selects a trace or rake of traces for display and the traces are computed and then drawn—in most cases almost instantaneously—much like a smoke wand in a wind tunnel.

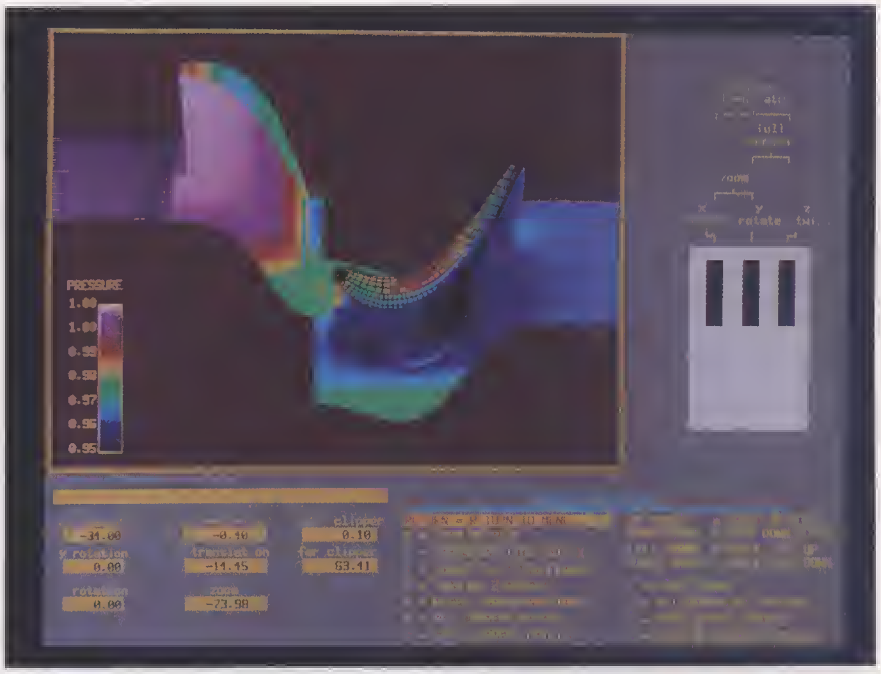

Figure 5. This is an example of Surf (Surface Modeler). The scene shows a computer-generated model of the rotor-stator interaction in a turbine engine. The objects are rendered partly as computational grids and partly as smooth-shaded solids, and the surface is color-mapped to show pressure. The Surf menus and mouse diagram are also visible.

Discussion

Of the three visualization methods (postprocessing, tracking, and steering) discussed in this article, postprocessing is by far the most common. Current supercomputer and workstation performance make this the most practical method for viewing solutions of computational aerodynamic solutions. Postprocessing techniques include

(1) dynamic, interactive viewing on the workstation,
(2) recording and playback on videodisc and then to tape, and
(3) recording and playback on 16-mm film.

These recording techniques have greatly improved the ability of scientists at NASA Ames to conduct fluid dynamics research, although these techniques necessarily mean a loss of interactivity, take a long time, and, for video, mean a loss of spatial and color resolution.

With direct viewing on the workstation, the capability to interactively manipulate the viewing position and the animation sequence was found to be very effective in providing a quicker and more complete understanding of the flow-field solutions. This capability is lost if the scenes are so complex that they must be recorded for playback. A solution for this problem is to increase the performance of the workstation.

The display creation speeds of workstations are projected to increase an order of magnitude over the next year. This will permit many complex scenes that now must be recorded for satisfactory motion analysis. There will still be scenes too complex to view with adequate rates of motion on the new workstations; recording will still be required for these scenes. In addition, recordings are required for group presentations. Therefore, it is important to improve the recording techniques also.

Recording on the Abekas videodisc or Dunn film recorder requires a much longer time than simple viewing on the workstation—a typical recording time is 30 to 60 minutes for every one minute of playback time. The Abekas video system records each frame essentially instantaneously, so the length of recording time is determined by the time it takes to render each frame, which is determined by the rendering techniques being used for the simulation (1 second to 2 minutes).

Video recording could be substantially improved with the addition of real-time digital video output from a workstation frame buffer. Certain digital video component manufacturers (Sony, Ampex, etc.) are already standardizing on the D2 composite digital video format and, although many workstation manufacturers have discussed such an option, availability is limited at the time of this writing; the limited availability is only in the D1 component digital video format. Not only would this option eliminate the need for much of the outboard equipment necessary for video recording, it could potentially improve the video quality by eliminating numerical sampling errors going from digital to analog and back to digital again.

While there will always be a demand for presentation videos, partially reducing the dependence on these recording techniques would require ultimate performance in a computer-graphics workstation. A spatial resolution of 1,280 × 1,024 requires 100,000 polygons per second updated at 10 to 12 frames per second for baseline performance (with hidden surfaces removed, antialiasing, and interactive, advanced lighting models). Workstations that approach this level of performance are discussed in this article. Other possibilities include fast frame-buffer configurations utilizing a very high speed network interface to a supercomputer (100 megabytes per second) or an RGB digital video system (although this would not be interactive).

As further advances are made in supercomputing, parallel architectures, networks, and workstation graphical performance, we predict the development of more and more software environments where tracking and steering techniques are employed. State-of-the-art resources allow for only minimal examples of these types of scientific visualization of computational aerodynamics.

The high-resolution, high-performance 3D graphical workstation combined with specially developed display and animation software has provided the scientists conducting fluid-flow simulations with good tools for analyzing flow-field solutions obtained from supercomputers. Tracking and steering, the two other methods discussed that allow interaction with the solutions as they evolve, are desirable but not yet practical, given the current state of computer technology. Postprocessing, the creation of graphics after the solutions are completed, is currently the standard method for visualization in computational aerodynamics.

A video system or 16-mm film recorder, plus the controlling animation software, is needed in addition to the workstation for very complex scenes that cannot be created rapidly enough now to yield satisfactory dynamics on the workstation alone.

Scientific visualization turns computations into pictures. Without the visualizations, the aerodynamic simulations are an unintelligible mass of numbers; with the visualizations, the simulations provide usable, viewable, transportable information that can be of great scientific value when analyzing the results of computational fluid simulations. □

References

1. P. Buning et al., "Flow Visualization of CFD Using Graphics Workstations," AIAA 87-1180, *Proc. 8th Computational Fluid Dynamics Conf.,* June 9-11, 1987.
2. J. Foley and A. Van Dam, *Fundamentals of Interactive Computer Graphics,* Addison-Wesley, 1982.
3. T. Lasinski, "Second Generation Graphics Workstations, Request For Proposals," RFP2-33491, Jan. 1989.
4. P. Buning and J. Steger, "Graphics and Flow Visualization in Computational Fluid Dynamics," AIAA-85-1507-CP, *Proc. AIAA 7th Computational Fluid Dynamics Conf.,* July 15-17, 1985.
5. D. Choi and C. Levit, "An Implementation of a Distributed Interactive Graphics System for a Supercomputer Environment," *Proc. 2nd Int'l Conf. on Supercomputing (ICS),* Mar. 1987.
6. S. Rogers, P. Buning, and F. Merritt, "Distributed Interactive Graphics Applications in Computational Fluid Dynamics," *Int'l J. Supercomputing Applications,* Vol. 1, No. 4, Winter 1987.

Gordon V. Bancroft is task manager of a Sterling Software group under contract to NASA Ames Research Center to develop tools for visualization of computational aerodynamics. The group is responsible for developing and supporting such tools as GAS (Graphics Animation System), Surf (Surface Modeler), and FAST (Flow Analysis Software Toolkit). The group is also responsible for supporting videotape and film production of NASA's Computational Fluid Dynamicists (CFD) animation results for the Numerical Aerodynamic Simulation facility at NASA Ames.

Readers can contact Bancroft at NASA Ames Research Center Workstation Applications Office—Code RFW, MS 258-2, Moffett Field, CA 94035.

Todd Plessel is a programmer/analyst with Sterling Software, a contractor at NASA Ames Research Center. Since joining the Workstation Applications Center in 1986, he has worked on the Graphics Animation System (GAS) and authored Surf, a surface modeling program. He is currently the lead designer/programmer for the Flow Analysis Software Toolkit (FAST) program, a software environment for NASA's Computational Fluid Dynamicists (CFD).

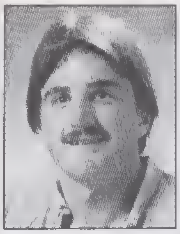

Fergus Merritt is working on continuing enhancements to the Graphics Animation System (GAS) and other interactive Computational Fluid Dynamicists (CFD) graphics programs at the NASA Ames Research Center. He has been with NASA Ames since 1984 and was one of the original members of the Workstation Applications Office. One of his first assignments was to design and develop the GAS package.

Merritt received his BS degree in computer science from the University of California at Berkeley in 1984.

Pamela P. Walatka, a technical writer for Sterling Software at the NASA Ames Research Center, writes manuals for computer graphics applications and libraries. She coauthored the *Plot3D User's Manual* and is currently working on the *Panel Library Programmer's Manual.*

Walatka received her BA from the University of California at Berkeley.

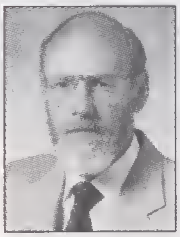

Val Watson is a scientist at the NASA Ames Research Center conducting research in computational fluid dynamics and computer graphics. He is currently chief of the Workstation Applications Office, which is developing programs of visualization of flow fields using high-performance graphics workstations.

Watson received a PhD in aeronautics and astronautics from Stanford University in 1969. He is a member of the IEEE, the ACM, the American Institute of Aeronautics and Astronautics, and the American Society for Engineering Education.

Visualization of Experimental Data at the Naval Research Laboratory

Lawrence J. Rosenblum, Naval Research Laboratory

The Naval Research Laboratory began operations in 1923 when it was created in response to Thomas Edison's suggestion that "The government should maintain a great research laboratory ... In this could be developed ... all the techniques of military and naval progression without any vast expense."

The original two NRL divisions, Radio and Sound, pioneered in the fields of high-frequency radio and underwater sound propagation, producing communications equipment, direction-finding devices, sonar sets, and the first practical radar equipment built in the United States.

New divisions were created, and today NRL's mission is to conduct a broadly based multidisciplinary program of scientific research and advanced technological development directed towards new and improved materials, equipment, techniques, and related operational procedures for the US Navy. Fulfilling this mission has resulted in research programs covering a wide spectrum of scientific areas from astronomy and satellites to the deep ocean; from studies of the structure of materials to radar, electronics, and computer science. The ability to visualize scientific data, that is, produce data products in graphical formats that enable the human eye/brain to perceive and infer patterns, is playing an increasingly important role in helping NRL achieve its goal of developing and transitioning technological innovation to the fleet and to society as a whole.

NRL provides its researchers with a Cray and a Connection Machine for large-scale computation, as well as numerous midsize computer facilities and desktop computers. The problems associated with providing appropriate graphics for numerical simulations on supercomputers at NRL are similar to those discussed elsewhere in this issue of *Computer* for other large laboratories. However, much of the science at NRL (and elsewhere) is experimental and, with measuring devices producing ever higher data rates, the need to visualize experimental data is now inherent to the data analysis process for many disciplines.

In computational science, the problem is to create appropriate models. Once this is achieved, visualization can typically be directly performed on the numerical simulation data. On the other hand, for experimental data, the parameters of interest must be extracted from the data sets. For example, in physical oceanography, temperature is measured, but a quantity called *activity* (see the section entitled "Physical oceanography") is calculated and visualized. Similarly, in molecular modeling, diffraction intensities are measured and electron densities are calculated and used to model the structure. This article will focus on applications of visualization to experimental data, discuss the methods used to derive parameters of interest, and show the role of visualization in interpreting the data.

Physical oceanography

Much of the earth's surface is covered by water. The need to model the global climate places increased importance on understanding ocean processes because of their strong coupling to the atmosphere. Similarly, naval systems for surface operations and antisubmarine warfare require detailed knowledge of the ocean environment. The ocean is a particularly difficult medium to study; it is a 3D space that varies temporally and is filled with inhomogeneities of all scales (from microbubbles to whales) while being subject to a wide variety of forcing functions, including thermodynamic, seismic, atmospheric, inertial, and tidal.

To study ocean processes that range in length scales from meters to kilometers, NRL scientists tow two primary devices through the ocean. A vertical array of closely spaced thermistors provides temperature measurements of the upper ocean, while a Doppler shear profiler uses sonar to measure the magnitude and direction of subsurface vertical shear of the current (i.e., differences in the currents at different depths). After processing, these data provide important clues about physical processes in the ocean.

Figure 1, which shows measured shear magnitude and direction as obtained from Doppler profiler data, demonstrates the role of visualization. Both magnitude and direction are seen to be horizontally layered. The strong shear magnitude near the surface (red bands in the upper panel) suggest forcing by variations in the trade winds. In the lower panel, the direction can be seen to cycle five times, indicating five cycles of a ''near-inertial'' internal wave. The sense of color (i.e., direction) progression implies that the energy in these waves is moving downward. This is consistent with the input of energy at the surface by the wind. The ability to visualize field-collected data in this manner not only provides knowledge directly, but also assists the oceanographer in prescreening data for further analysis. Figure 1, shown without axes and with certain processing anomalies visible, also illustrates the concept of ''peer graphics.'' Like most scientific graphics, this figure was not made for publication but rather for previewing the data and for discussion with colleagues. The figure was created with software that allows the quick display of any pair of 2D data sets and gives the user several choices of color selection for each set.

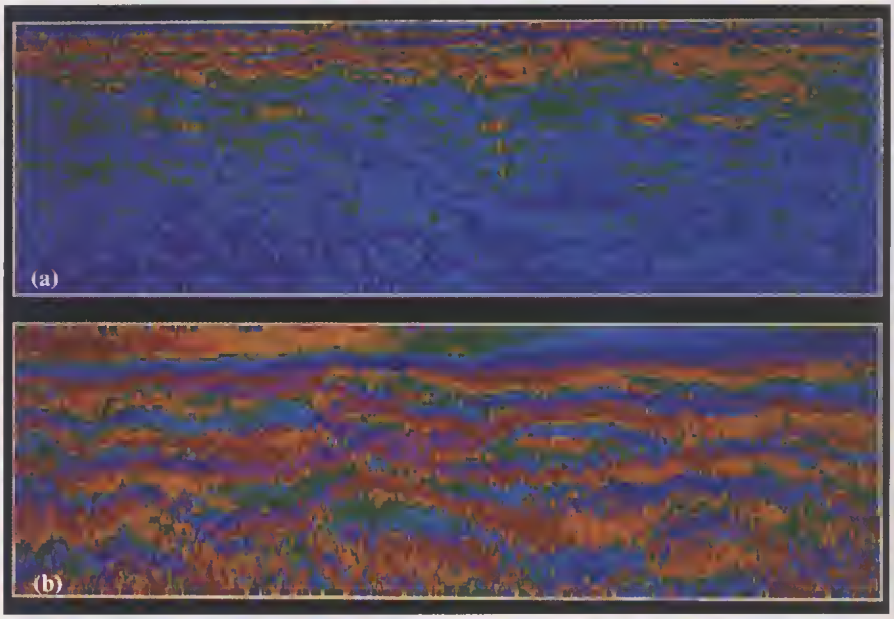

Figure 1. Ocean currents: a vertical slice through the upper 300 meters of ocean east of Barbados, for 130 kilometers of horizontal tow. Panel (a) shows the magnitude of the vertical shear with intensity increasing across the color scale from blue to red and violet. Panel (b) shows shear direction, color coded from 0 to 2 π radians. This figure is typical of visualization tools used for a ''quick look'' at oceanographic data sets.

The use of temperature data is more complex, because our primary interest lies in activity; that is, the local variability of temperature. A fast Fourier transform-based algorithm[1] is used to extract activity values from thermistor ''chain'' measurements. The resulting activity data are interpolated, color coded, median filtered, and displayed (see the top panel of Figure 2). To compare fine-structure activity with other parameters, the activity field is edge detected. The resulting outlines of high activity ''patches'' are overlaid on fields of physical parameters conjectured to relate to the generation of enhanced activity (see the five lower panels of Figure 2). This visualization technique makes obvious the strong correlation of enhanced activity with high shear and low Richardson number (a theoretical predictor of turbulence), suggesting that shear instabilities are creating the enhanced fine-structure activity.[2]

Electronic warfare

Neutralizing an incoming missile is fundamental to the survival of the fleet in any modern-day naval encounter. The Central Target Simulator facility, a laboratory containing computer-controlled simulators, was created to develop protective countermeasures for ships and aircraft from radio frequency guided missiles. Figure 3a shows a simulator that allows both the missile hardware RF section and the computer section to be simulated as if the missile were actually flying against a target. Using the simulator, dozens of simulations can be performed quickly to test the vulnerability of the fleet to missile attack.

The yellow circles on the wall of the simulator's anechoic chamber are the conical sections of microwave antennas driven from an RF switching network. The network is digitally controlled to allow transmission of RF signals that create a tactical environment simulating the radar and guidance systems of a missile. A multiprocessor architecture is used to move the targets over this antenna array, creating an illusion of motion. Various engagement scenarios are created to test the missile, the target's defense, or a missile's susceptibility to electronic countermeasures (ECM).

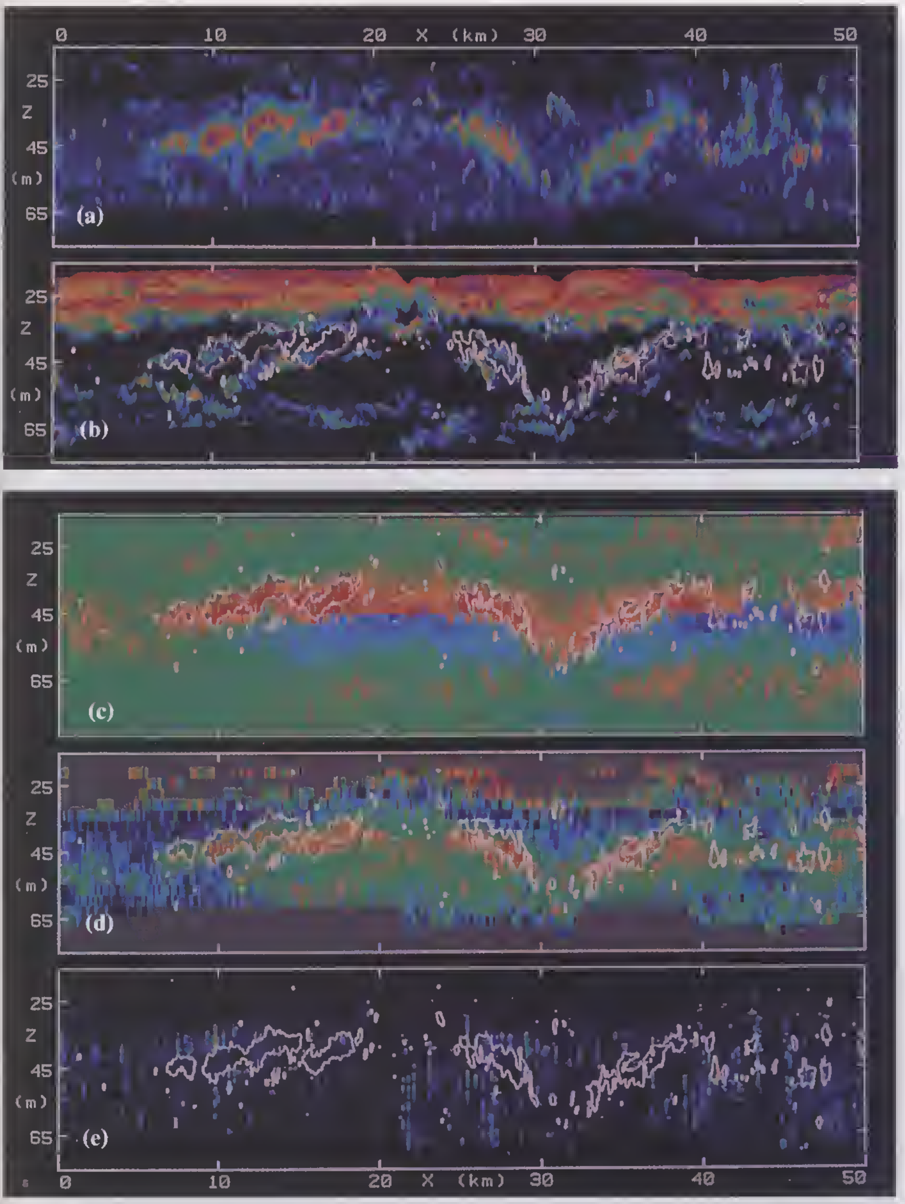

Figure 2. From a Sargasso Sea experiment where a high activity "patch" was located and repeatedly transected. Panel (a) shows activity, with yellow to red indicating high activity regions. The remaining panels show high activity contours overlaid on (from the top) (b) vertical temperature gradient, (c) shear, (d) Richardson number, and (e) internal wave potential energy. Yellow to red indicates low or inverted temperature gradient, high shear, low Richardson number, and high IWPE.

Visualization tools are necessary to portray the complex data contained in a simulation in real time. Patterns not readily observable in simulation data become apparent when viewed pictorially. Animations of simulation test results are also used to confirm that the correct real-world environment is being simulated and that the hardware is being properly exercised, as well as to aid in analysis of missile and ECM effects.

Figure 3b is a frame from a computer-generated animation of a missile/ship engagement. The generation of the ship, the chaff cloud, the missile, and the missile beam use classical polygon rendering/animation techniques. The moving background of sea and clouds are generated using an implementation of fractal mathematics.[3] The realistic ocean and clouds give a sense of motion to the visualization, increasing the viewer's comprehension of the spatial dynamics. In addition to creating a more realistic looking scene, fractal methods are computationally less costly than polygon techniques.

Molecular modeling

The accurate determination of the atomic arrangements of substances in various states allows the scientist to understand the relationship between the structural and the physical, chemical, and biological properties of a substance. Products of molecular modeling range from new, improved antibiotics to safer, more powerful propellants for solid-fuel rockets. The need to visualize molecules predates computer graphics; chemists have long "visualized" by using a Tinkertoy-like construction set. Molecular modeling is one of the oldest areas of computer-based scientific visualization, and one where, hand-in-hand, both fields have helped the other progress further.

The mathematics used to help solve molecular structures were developed at NRL. Jerome Karle, head of NRL's Laboratory for the Structure of Matter, and Herbert Hauptman, a former LSM member, received the 1985 Nobel Prize in Chemistry for their use of the mathematics of diffraction physics to show that enough information could be deduced from diffraction patterns to permit the direct solution of crystal structures—the "direct method" for molecular determination via x-ray crystallography. Isabella Karle of the LSM has also made major contributions to the development of the methods and procedures for structure determination, including the molecular structures illustrated in Figures 4 and 5.

The output of a molecular structure determination by x-ray crystallography is a numerical set of atomic positions. Visualization enables the mind to grasp the structure and its implications. Two common visual representations are the space-filling model and the ball-and-stick model. The space-filling model is useful for analyzing potential sites for molecular interactions. With such knowledge, scientists can create new, useful molecules. Figure 4, from a study of synthetic antibiotics, illustrates the space-filling model.[4]

The ball-and-stick model is superior for visualizing the overall structural configuration of a molecule. It shows the bonds between atoms and enables visualization of the angles between the bonds and of the distances between the atomic centers. By imagining the breaking of given bonds, the formation of new bonds, and a flexing of the structure, the mechanisms of chemical reactions can be understood. Figure 5 is one of a stereo pair that represents the product of a photochemical reaction with the two components of the molecule joined together with a central cage structure.[5] With practice, scientists can view a stereo pair (without stereo viewing devices) and "see" the molecule in three dimensions. Computer-based molecular modeling offers the user many advantages, including color selection, zooming in on features of interest, user-controlled interactive movement, easy generation of molecules with a very large number of atoms, and matching of multiple molecules for determining possible sites for molecular interactions. The molecules were generated on a VAX 8650 using custom software developed in-house.

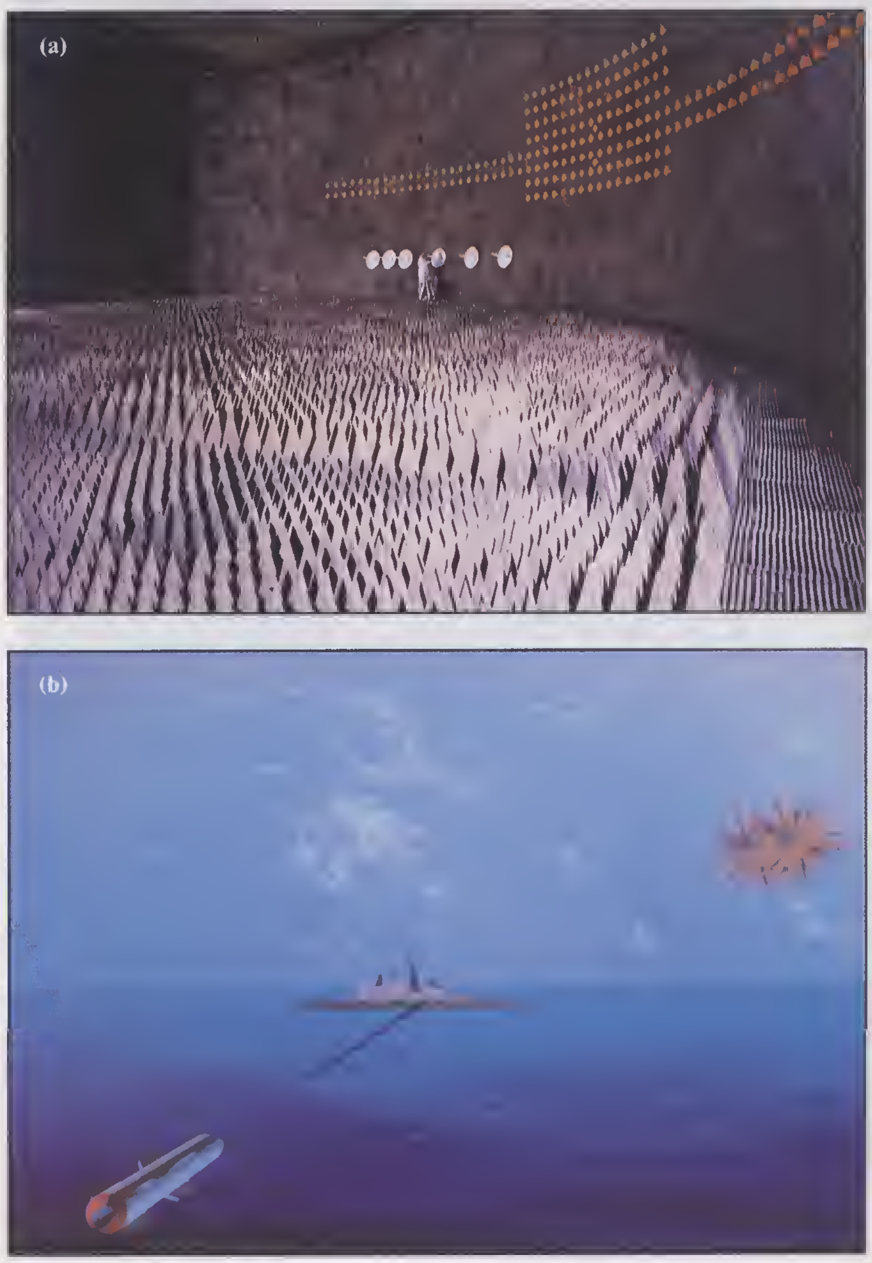

Figure 3. Electronic countermeasures: The ECM laboratory is shown in (a). The size can be seen by comparison with the person standing at center rear. A frame from an ECM animation is shown in (b). The chaff cloud to the right of the ship is an attempt to mislead the missile.

Underwater acoustics

Matched-field processing[6] is a new technique for processing ocean acoustic data measured by an array of hydrophones to produce an estimate of the location of sources of acoustic energy, including whales, volcanoes, and submarines. This method differs from source localization techniques used in other disciplines, for example, radar and astronomy, in that it attempts to exploit the complicated underwater acoustic environment to improve the accuracy of the source location estimate.

Central to the idea of matched-field processing is the premise that an ocean acoustic propagation model can reliably be used

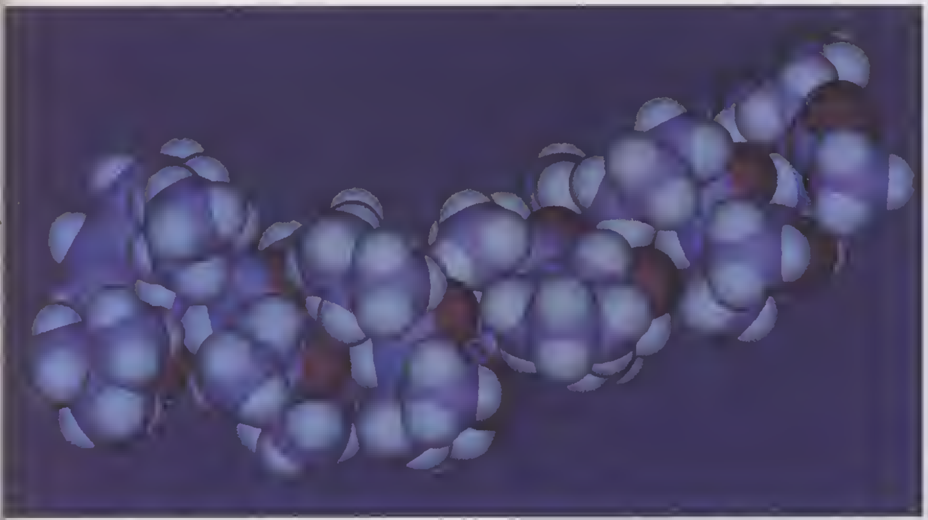

Figure 4. Space-filling model of a molecular structure used in the study of synthetic antibiotics. It represents a synthetic apolar analog of the membrane-active fungal peptide antibiotic zervamycin IIA. Atoms are color coded: hydrogen, white; carbon, violet; oxygen, red; and nitrogen, blue.

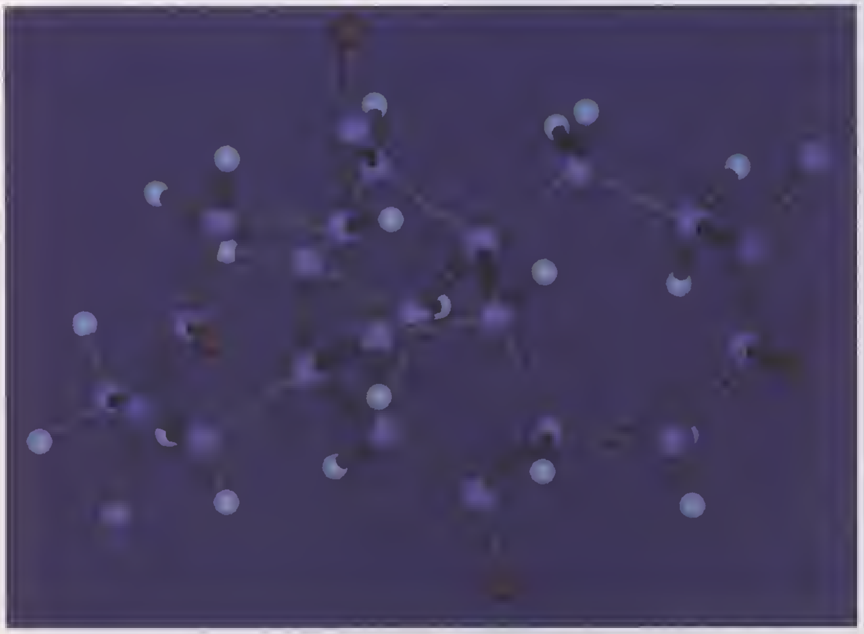

Figure 5. Ball-and-stick model of a product of a photochemical reaction, referred to as a photodimer from N-Chloroacetyltyramine. Atoms are color coded: hydrogen, white; carbon, violet; oxygen, red; and nitrogen, blue.

to compute a replica of the acoustic field measured by the hydrophone array. Using one of a number of proposed algorithms,[7] the computed replica field is matched with the measured acoustic field to produce a 3D ambiguity measure as a function of the unknown source location variables, for example, range, bearing, and depth. The best estimate of the source location is then obtained from the location of the peak of the ambiguity function. To succeed, both the specific ocean environment and the array configuration must generate an acoustic field sufficiently unique to produce an ambiguity function where the peak can be identified against the background. In addition, uncertainties in the locations of the hydrophone elements and the environment as well as modeling errors must not significantly degrade the identification.

Visualization plays an important role in the analysis of the ambiguity function. To develop understanding of the matched-field process, not only must the peak be identified, but the structure inherent within the 3D ambiguity function needs to be understood. In Figure 6a, the ambiguity surface volume has been generated using an isosurface method. The plane facing the reader represents a range-bearing slice for fixed depth; the top plane, range-depth for fixed bearing; and the side plane, depth-bearing for fixed range. In this example, the acoustic source is located in the center and two false target ''sidelobes'' can be seen at the left and right ends of the volume. The scientist develops understanding of the acoustic effects by interactively moving through the volume (for example, Figure 6b) in real time, using dials on a knob-box to control direction. Efforts are under way to apply volume rendering techniques[8] to this problem, both to provide even better visualization of data such as in Figure 6 and, in the future, to develop methods to understand dynamic (i.e., time varying) ambiguity functions.

Radar

Synthetic aperture radar (SAR) is a type of radar that constructs a detailed 2D image by utilizing the motion of the platform to synthesize the effect of a large antenna aperture.[9,10] The properties of the electromagnetic radiation for the 3-to-30-centimeter wavelengths used for SAR make it very useful for surveillance and remote sensing of the earth's surface. Day or night, SAR penetrates clouds, rain,

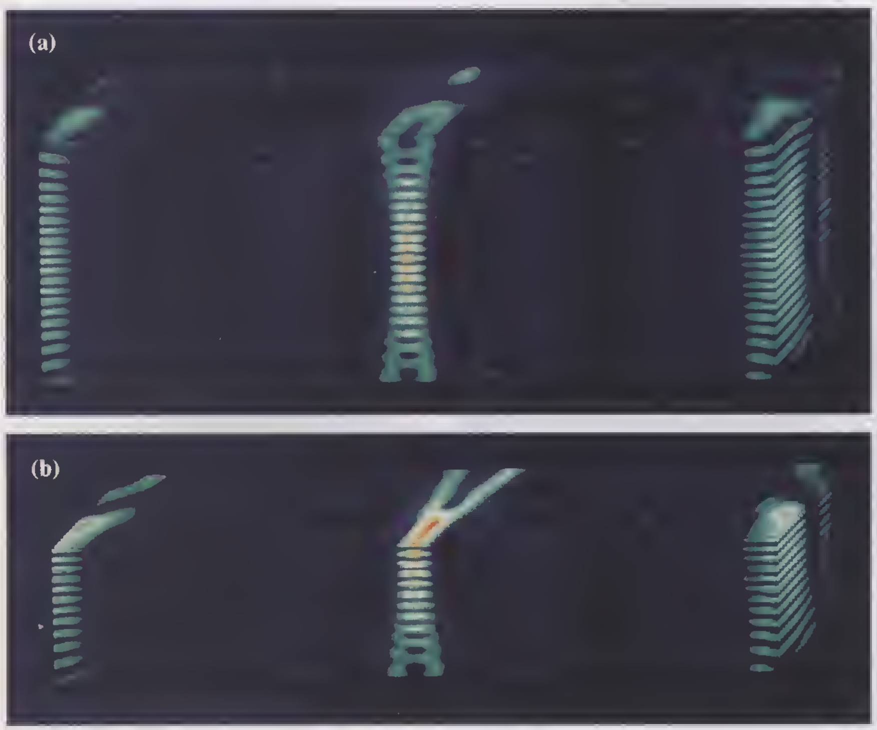

Figure 6. Isocontour rendering of an acoustic ambiguity function. Axes are range (horizontal), bearing (vertical), and depth (into the screen). The acoustic source is at the center of the volume (green-to-red). Focused acoustic energy can be seen at the left and right portions of the volume as well. The full ambiguity function is shown in (a), while (b) shows part of the volume after interactively slicing into it.

and, to some degree, foliage, yielding a clear picture of the surface under most weather and ground-cover conditions. Used over the ocean, SAR can provide indications of sea state, subsurface ocean processes, and even some bathymetric (ocean floor) features.

Typically, SAR is flown on an aircraft or satellite and a scene is illuminated by short frequency-modulated radar pulses at rates of 1,000 to 2,000 pulses per second. Typical data recording rates are between 10 and 90 megabits per second. The antenna beamwidth defines the illuminated area. Complex processing is required to reconstruct the SAR image from the received signals.

To obtain range resolution, a reference matched FM filter compresses the return pulses after they have scattered from the surface. The azimuth (along track) resolution is formed synthetically as the SAR system moves past a scene. Echo amplitude is sampled and phase history reconstructed during the time period between consecutive pulses. This forms the synthetic aperture that can contain several thousand points along the azimuth direction. A matched Doppler filter is correlated with the reconstructed phase delay to compress the signal in azimuth. Each pixel is then formed, until a complete image emerges. Corrections are made for effects such as the earth's rotation beneath the SAR platform that would otherwise distort the image.

Visualization is essential to interpret the large quantities of SAR data; rich details can easily be seen in the images. Figure 7, acquired from a 1984 experiment on NASA's Challenger space shuttle (Shuttle Imaging Radar, SIR-B), shows a portion of the Atlantic Ocean, 45 kilometers in azimuth by 34 kilometers in range, some 135 kilometers south of Long Island.

Although SAR penetrates the ocean's surface only to a depth of about one inch, many features of the ocean can be seen.[9] One feature is the continental shelf (upper right quadrant) in which the image brightens over the rapid drop-off to the deep ocean floor (lower half) as depth changes abruptly from 100 meters for the continental shelf to over 700 meters in the Hudson Canyon.

SAR images can be used to identify oceanographic effects caused by wind, wave, and current interactions. Periodic waves can be seen in the image, especially in the center right portion. These are subsurface internal gravity waves, triggered by currents flowing across the edge of the continental shelf.

The above examples illustrate how visualization is helping NRL investigators understand experimental data obtained from complex measurements, taken in the laboratory, in the air, and at sea. The techniques—signal processing, image processing, animation, polygon-rendering, fractals, and volumetric methods—and results shown illustrate the diverse methods required for experimental data from different disciplines. The multidisciplinary examples provide a sampling of how visualization is enhancing scientific knowledge. The diversity also demonstrates the difficulty in developing generalized visualization tools for experimental data.

However, it is reasonable to expect that mergers of certain techniques will become common. In particular, volumetric methods, used primarily to display "real" objects (e.g., medical and stress analysis applications), are increasingly being used to display spaces consisting of arbitrary, but related, parameters. Thus, signal processing, image processing, and volume visualization will combine to permit effective interpretation of the structure of multidimensional spaces. The need to perform time-varying volume rendering is becoming important.

Primitive forms of computer-based visualization have been with us since the earliest computers, but the focus on visualization as a discipline is new. As computers and measurement devices become increasingly powerful, the role of visualization in the discovery process will become even more essential. New visualization techniques will be required, as will structures for their effective application. As these develop, visualization will provide even deeper insight and understanding into the real world around us. ☐

Figure 7. Synthetic aperture radar image of the Atlantic Ocean south of Long Island. The 1984 image was taken at 4:03 a.m. EST from the Challenger space shuttle at a height of 218 kilometers. The image is color-coded, with dark representing low intensity; blue, middle intensity; and white, high intensity regions of backscattered radiation. Both ocean floor features and subsurface wave interactions can be seen.

Acknowledgments

The work discussed in this article represents the efforts of many people: physical oceanography, George Marmorino, Cliff Trump, and the author; electronic warfare, Bob Normoyle (Quest Corp.) and Tony Ricci; molecular modeling, Ronald M. Brown from data developed by I.L. Karle and others; matched-field visualization, Dick Heitmeyer and the author; and SAR radar, NRL's Digital Image Processing Laboratory. I thank all the above for their assistance in developing this article as well as Jim Lucas et al. in NRL's Technical Information Division for their assistance in producing the supplemental video.

References

1. L.J. Rosenblum, "Visualizing Oceanographic Data," *IEEE Computer Graphics and Applications,* May 1989, pp. 14-19.
2. G.O. Marmorino, L.J. Rosenblum, and C.L. Trump, "Fine-Scale Temperature Variability: The Influence of Near-Inertial Waves," *J. Geophysical Research,* Vol. 92, No. C12, Nov. 1987, pp. 13,049-13,061 and 13,215-13,219.
3. *The Science of Fractal Images,* H.O. Peitgen and D. Saupe, eds., Springer-Verlag, New York, 1988.
4. I.L. Karle et al., "Conformation of a 16-Residue Zervamicin IIA Analog Peptide Containing Three Different Structured Features: z10-helix, a-helix and b-blend ribbon," *Proc. Nat'l Academy of Science,* Vol. 84, Aug. 1987, pp. 5,087-5,091.
5. D.S. Jones and I.L. Karle, "The Crystal and Molecular Structures of Two Photodimers from N-Chloroacetyltyramine," *Acta Crystallographica,* Vol. B30, part 3, Mar. 1974, pp. 617-623.
6. A.B. Baggeroer, W.A. Kuperman, and H. Schmidt, "Matched-Field Processing: Source Localization in Correlated Noise as an Optimum Parameter Estimation Problem," *J. Acoustic Society America,* Vol. 83(2), 1988, pp. 571-587.
7. M.B. Porter, R.L. Dicus, and R.G. Fizell, "Simulation of Matched-Field Processing in a Deep-Water Pacific Environment," *IEEE J. Oceanic Engineering,* Vol. OE-12, No. 1, 1987, pp. 173-181.
8. K.A. Frenkel, "Volume Rendering," *Comm ACM,* Vol. 32, No. 4, Apr. 1989, pp. 426-435.
9. R.O. Harger, *Synthetic Aperture Radar Systems,* Academic Press, N.Y., 1970.
10. B.L. Huneycutt, "Spaceborne Imaging Radar-C Instrument," *IEEE Trans. Geoscientific and Remote Sensing,* Vol. 27, No. 2, Mar. 1989, pp. 164-169.

See the final page of the introduction to this section for Rosenblum's biography, photograph, and address.

Advance Program

Eighth Symposium on Reliable Distributed Systems

Sponsors
IEEE Computer Society
TC on Distributed Processing

In cooperation with
TC on Fault-Tolerant Computing
Seattle University

October 10-12, 1989, Stouffer-Madison Hotel, Seattle, Washington

Honorary Chair
Dr. Barry Boehm
TRW Inc.

Symposium Chair
Professor Jane W.S. Liu
University of Illinois

Program Chair
Dr. Raif M. Yanney
TRW Inc.

Program Committee
Professor Bharat K. Bhargava
Purdue University

Professor Susan Davidson
University of Pennsylvania

Professor John P. Hayes
University of Michigan

Professor John Kelly
University of California, Santa Barbara

Professor K.H. (Kane) Kim
University of California, Irvine

Dr. John Knight
Software Productivity Consortium

Professor Herman Kopetz
Vienna, Austria

Mr. Thomas Lawrence
RADC

Professor Miroslaw Malek
University of Texas, Austin

Dr. Edgar Nett
Bonn, West Germany

Dr. Thomas Raeuchle
Honeywell, Inc.

Professor David Rennels
University of California, Los Angeles

Professor Richard Schlichting
University of Arizona, Tucson

Dr. Luca Simoncini
IEI-CNR, Pisa, Italy

Dr. Basil Smith
IBM TJ Watson Research Center

Dr. Wing N. Toy
AT&T Bell Laboratories

Dr. Anand Tripathi
University of Minnesota

Tutorial Chair
Professor Kwei Jai Lin
University of Illinois

Treasurer
Professor Pui Ng
University of Illinois

Local Arrangements
Professor Ihsin Phillips
Seattle University

Publications
Dr. David Cohen
IEX Corporation

Steering Committee
Professor Ming T. Liu
Ohio State University

Tuesday, October 10, 1989

8:00 am – 9:00 am	Registration
9:00 am – 5:00 pm	Tutorial
6:00 pm – 8:00 pm	Reception

Wednesday, October 11, 1989

7:30 am – 8:30 am	Registration
8:30 am – 10:00 am	Plenary Session Opening Remarks: Jane W.S. Liu, University of Illinois Raif M. Yanney, TRW Inc. Keynote Speech: TBD
10:00 am – 10:30 am	Break
Session 1: 10:30 am – 12:00 noon	Recovery in Distribution Systems/Chair: Susan B. Davidson, University of Pennsylvania **"Using Checkpoints to Localize the Effects of Faults in Distributed Systems"** Mustaque Ahamad and Luke Lin, Georgia Institute of Technology **"A Low Overhead Checkpointing and Rollback Recovery Scheme for Distributed Systems"** Zhijun Tong, Richard Y. Kain and W.T. Tsai, University of Minnesota **"Recovery-Management in the RelaX Distributed Transaction Layer"** R. Schumann, R. Kroger, M. Mock and E. Nett, GMD, West Germany
12:00 noon – 1:30 pm	Lunch
Session 2: 1:30 pm – 3:00 pm	Replication/Chair: Thomas Raeuchle, Honeywell, Inc. **"Consistent Replicated Transactions: A Highly Reliable Program Execution Environment"** Shepherd S.B. Shi and Geneva G. Belford, University of Illinois at Urbana-Champaign **"Implementing Fault-Tolerant Replicated Objects Using Psync"** Shivakant Mishra, Larry L. Peterson and Richard D. Schlichting, University of Arizona **"Recovering from Process Failures in the Time Warp Mechanism"** Jonathan Agre, Rockwell International Corp., and Divyakant Agrawal, University of California, Santa Barbara
3:00 pm – 3:30 pm	Break
Session 3: 3:30 pm – 5:00 pm	Reliable Communication/Chair: Anand Tripathi, University of Minnesota **"Self-Stabilization of the Alternating-Bit Protocol"** Yehuda Afek and Geoffrey M. Brown, Cornell University **"A Token-Based Protocol for Reliable, Ordered Multicast Communication"** B. Rajagopalan and P.K. McKinley, University of Illinois **"An Efficient Kernel-Level Dependable Multicast Protocol for Distributed Systems"** M. Stella Atkins, Garnik Haftevani and Wo Shun Luk, Simon Fraser University
5:00 pm – 7:00 pm	Reception

Thursday, October 12, 1989

7:30 am – 8:30 am	Registration
Session 4: 8:30 am – 10:00 am	Performance Analysis/Chair: Tom Lawrence, RADC **"Workload Analysis for Performance Study of Distributed Deadlock Detection Algorithms"** Kyung O. Min and Geneva G. Belford, University of Illinois at Urbana-Champaign **"Performance Comparison of Concurrency Control Protocols for Transaction Processing Systems with Regional Locality"** Bruno Ciciani, Daniel M. Dias and Philip S. Yu, IBM TJ Watson Research Center **"Implementing Atomic Rendezvous within a Transactional Framework"** J.P. Banatre, M. Banatre and C. Morin, IRISA, France
10:00 am – 10:30 am	Break
Session 5: 10:30 am –12:00 noon	Evaluation and Modeling/Chair: Dave Rennels, University of California, Los Angeles **"Availability Analysis of the Primary Site Approach for Fault Tolerance"** Yennun Huang and Pankaj Jalote, University of Maryland **"Bounded Approximate Reliability Models for Distributed Systems"** Joanne Bechta Dugan, Mark Boyd and Nitin Mittal, Duke University, and Malathi Veeraraghavan, AT&T Bell Laboratories **"Reliability and Scalability of Large-Scale Gracefully Degradable Systems"** W. Najjar and J-L. Gaudiot, University of Southern California
12:00 noon – 1:30 pm	Lunch
Session 6: 1:30 pm – 2:30 pm	Network Architecture/Chair: Wing Toy, AT&T Bell Laboratories **"Ring Network Reliability – The Probability That All Operative Nodes Can Communicate"** Nanchang Lin and Charles B. Silio, Jr., University of Maryland **"Decentralized Loopback Reconfiguration of a Bidirectional Ring LAN"** Jernej Polajnar, Duro Mastilovic and Blagoja Pardulski, Boris Kidrit Institute of Nuclear Sciences, Yugoslavia
2:30 pm – 3:00 pm	Break
Session 7: 3:00 pm – 4:00 pm	Simulation and Testing/Chair: John Kelly, University of California, Santa Barbara **"Testing Reliable Distributed Applications through Simulated Events"** Travis L. Winfrey and Gail E. Kaiser, Columbia University **"Distributed Processing Test Simulator"** Kenneth R. Leeper, Boeing Advanced Systems, and David W. Lyness, Tau Corporation

Eighth Symposium on Reliable Distributed Systems

Tutorial: Fault Tolerance and Evaluation Techniques for Real-Time Systems

Lecturers: Kang G. Shin, Department of Electrical Engineering and Computer Science, University of Michigan

C.M. Krishna, Department of Electrical and Computer Engineering, University of Massachusetts

Audience: Engineers and scientists who want an overview of the critical issues in fault-tolerant, real-time systems.

Abstract: Real-time computing systems which support critical applications, such as flight control, robotics, and machine vision, must satisfy certain timing and reliability constraints. This tutorial will survey the means currently available to accomplish this objective.

An effective performance measure to guide the design process of a system must convey concisely all the attributes of the system that impact its value for a particular application; traditional measures such as throughput and availability are inappropriate for real-time systems. Other performance measures, such as performability and cost functions, which are specifically developed for real-time systems and are application-sensitive, will be discussed. New analytical models to estimate the performance of real-time systems are emerging; they will be surveyed and their use will be discussed and illustrated with practical examples.

The second half of the tutorial will deal with techniques in fault-tolerant operating systems. Examples of these techniques include: schemes that maintain tight processor synchronization in the presence of arbitrary (or Byzantine) faults, task scheduling schemes that allow quick recovery from processor failure, and communication protocols appropriate for time-critical, distributed applications.

Biographies:

Kang G. Shin is a professor in the Department of Electrical Engineering and Computer Science at the University of Michigan. He has authored or co-authored over 140 papers in the areas of fault-tolerant, real-time computing; computer architecture, and robotics and automation. Founder of the university Real-Time Laboratory, he and his students are currently building a multiprocessor system called HARTS to validate various architectures and analytic results in real-time computing. He has been a visiting scientist at the U.S. Air Force Flight Dynamics Laboratory and at Bell Laboratories. His research in distributed real-time systems has been supported by NASA, the National Science Foundation, the Office of Naval Research, the Air Force Office of Scientific Research, and various corporations.

Dr. Shin was program chair for the 1986 IEEE Real-Time Systems Symposium (RTSS) and general chair of the 1987 RTSS. He also was guest editor of the August 1987 special issue of the *IEEE Transactions on Computers on Real-Time Systems.*

C.M. Krishna is an assistant professor in the Department of Electrical and Computer Engineering at the University of Massachusetts. He has conducted research on real-time computing and distributed systems, and co-authored more than 25 papers in these areas. He is currently establishing a real-time laboratory at the university. He has been a visiting scientist at the IBM Thomas J. Watson Research Center. His research has been supported by the National Science Foundation and the Office of Naval Research.

In 1986 Dr. Krishna was on the program committees of both the International Conference on Distributed Computing and the RTSS.

Symposium Registration Form

Complete and return this form with your check payable to *Reliable Distributed Systems Symposium* to: Professor Jane Liu
Department of Computer Science
University of Illinois at Urbana-Champaign
1304 West Springfield Avenue
Urbana, IL 61801
(217) 333-0135

Advance Registration (prior to September 8, 1989)

	IEEE/CS Member	Non-Member	Full-time Student
Tutorial	$170	$210	No Discount
Symposium	$210	$260	$50

Late Registration (after September 8, 1989)

	IEEE/CS Member	Non-Member	Full-time Student
Tutorial	$210	$260	No Discount
Symposium	$250	$310	$50

Total Amount Enclosed: ______

IEEE/CS Membership Number: ______

Name: ______

Organization: ______

Street Address: ______

City/State/Zip: ______

Country: ______

Hotel Registration Form for Reliable Distributed Systems Symposium

Hotel Address: Stouffer-Madison Hotel
515 Madison Street
Seattle, WA 98104
Tel: (206) 583-0300

Rates per day: Single/Double $105

A block of rooms has been reserved for the participants. To guarantee your reservation, it must be made directly with hotel prior to September 20, 1989.

Name: ______

Street Address: ______

City/State/Zip: ______

Single or Double Occupancy: ______

Arrival Date ______ Departure Date ______

M/C Number ______ Expiration Date ______

Visa Number ______ Expiration Date ______

Personal checks and travelers checks are also accepted.

"Any clod can have the facts, but having opinions is an art."
Charles McCabe, San Francisco Chronicle

THE OPEN CHANNEL

Jim Haynes, Computer Center, US Santa Cruz, CA 95065
Bitnet, haynes@ucscc; Arpa, ucscc!haynes@ucbvax.berkeley,edu

Stop revising, start finishing

Dwayne Phillips, US Defense Department

It happened again the other day at work. Someone revised what his predecessor had done. There really wasn't anything wrong with the way his predecessor had done it. He knew his predecessor — had known him for years. He respected his predecessor's ability. His predecessor was good at his job, but he never quite grasped what the project required. So now he just had to make a few changes. The changes would set the project back a few months, but that was acceptable because now things would be done just right — just his way.

In my nine years on the job I have seen this happen daily — people revising what their predecessors did instead of finishing the task. The task might be a weekly report or a $100 million project, but the result is the same: wasted time, effort, and money.

The most common revision is in the weekly or monthly report. Someone at the working level spends five minutes writing a paragraph on his or her activities for the week. A supervisor reviews it, groups it with the other paragraphs, and passes it up to the next level — but not before making a few revisions. Without revisions, his or her job wouldn't be necessary. The revisions take half a day. A further supervisor does the same thing, but revisions at the next level take two days. The process continues up the line until it takes a week to write and revise the weekly report.

In *The Psychology of Computer Programming*, Gerald Weinberg relates a similar story about a software company that had to submit monthly reports to the government. The review and revision process took a month. The result was that the workers had to predict what they were going to do instead of reporting what they had done.

Weinberg wrote this over ten years ago. His book wasn't read by many people.

These examples might seem amusing, but revising is also prevalent in serious and costly tasks. The most common place this occurs is the Pentagon, where people usually move from job to job every two or three years.

Here's the scenario: Someone moves into a project. His or her predecessor started the project two years before and it is finally getting off the ground. The new person sees the need for a few revisions and starts revising. The contractor stops working. Everyone concentrates on writing, studying, revising, and negotiating the revisions. This takes six to nine months, delays the project, and raises the price tag. After two years, the reviser moves on to another job. The reviser's replacement comes in, immediately spots a few things that aren't quite right, and makes a few revisions. The contractor again stops work for nine months. This delays the project and raises the price tag some more.

Are you mature enough to put aside your desire to leave your mark on the world?

This process continues through several more revisers. Slowly — very slowly — the project moves toward completion. The cost has grown out of sight and the schedule has slipped years beyond reason.

Why do people have to revise, revise, revise? Why can't they just finish the task?

Ego is one big reason. Some people simply know better than their predecessors. They must make revisions to correct the flaws in their predecessor's design. What they are really doing is correcting their predecessor — not their predecessor's design.

Job security is another reason for revising. Weinberg also mentioned this, but he called it job insecurity. Some people are so insecure in their ability to keep a job that they do things, such as revising, to make themselves indispensable. This is also known as justifying your job. You are an engineer (or programmer, or manager, or architect); therefore, you must make revisions or they wouldn't need an engineer on the job.

Another reason for revising is to leave your mark on the project. Years later, you can point to something in the system and claim it. Without you, that particular revision wouldn't be there.

In *The Psychology of Everyday Things*, Donald Norman calls this ''worshipping false images.'' People lose sight of the project's goal and start to place more importance on other things. The reviser worships himself and his revisions. The schedule delays and cost overruns are sacrifices that are part of this worship.

In some ways, revising is a question of maturity. Are you mature enough to put aside your ego, job insecurity, and desire to leave your mark on the world? Are you mature enough to see something through to completion? Are you mature enough to put the good of the project ahead of yourself?

Remember the goal of the project. Why are we all working here? The purpose is to build a system for defense, or medicine, or automobiles — not to revise our own image into the system.

There are times we can justify revisions. An innovation is sure to come along during the life of a long-term project, but don't revise it into the project unless it has a dramatic improvement (100 or 1,000 percent) on the performance, efficiency, safety, etc., of the project. Hold the innovation. Publicize it for the next project, but don't revise. Keep the project on course and finish it. Once you finish the project, try to put the innovation in a follow-up project. Management might disapprove the follow-up, but the original will be finished and working.

Revising is one of the big reasons that space shuttles, airplanes, automobiles, and government services are so expensive and time consuming. It is one of the reasons America keeps slipping in the world market. We need to stop revising and start finishing.

Definition of RISC

Daniel Tabak, George Mason Univ.

We read statements from time to time in the professional literature that there is no definition of the concept of reduced instruction-set computing (RISC) that is acceptable to most professionals. In view of the fact that there are about 20 RISC-type systems manufacturers and numerous literature items accumulated so far, perhaps the time has come to establish some standards.

With this open letter, I call on the president of the IEEE Computer Society to appoint a committee that will recommend a standard definition of RISC. I volunteer to serve on this committee. As a first iteration, I propose the following definition:

> A reduced instruction-set computer system is a computing system satisfying the following properties:
>
> (1) Streamlined and uniform treatment of all instructions by maintaining
>   (a) single-cycle execution of all instructions, and
>   (b) single-word standard length of all instructions (word length should equal the bus data width to permit a single-bus-cycle instruction fetch)
>
> (2) Reduced menu of possibilities related to all instructions, specifically
>   (a) 128 or fewer judiciously selected instructions,
>   (b) four or fewer addressing modes, and
>   (c) four or fewer instruction formats
>
> (3) Memory access by load/store instructions only
>
> (4) All operations are register-to-register within the CPU, except load/store.
>
> (5) Hardwired control
>
> (6) A relatively large (at least 32 registers), general-purpose, user-accessible, CPU register file (to facilitate the fourth point above, among other things).

In addition, a number of features in most existing systems can (and, in some cases, should) be implemented in any system, complex instruction-set computing (CISC) or otherwise:

> (1) High-level-language support in the design (i.e. reduction of the semantic gap and availability of optimizing compilers). This is essential for all systems, and cannot be regarded as an exclusive part of the RISC definition.
>
> (2) Implementation of register windows within the CPU register file
>
> (3) Pipelining
>
> (4) Delayed branch

I encourage comments from interested readers. My address is Electrical Engineering Department, King Hall, George Mason University, Fairfax, VA 22030.

STANDARDS

Editor: Fletcher J. Buckley, GE Aerospace, MS: 145-3, Moorestown, NJ 08057-0927; phone (609) 866-6350; fax (609) 866-6439; Compmail+ f.buckley

Transnationalization of IEEE Standards

Fletcher J. Buckley, GE Aerospace

We in the IEEE have a split personality with regard to standards. We draw our standards participation from all over the world, yet we act as a captive US standards organization.

On one hand, we claim to be a transnational organization with members in many countries. More than 21 percent of IEEE's 304,000 members live outside the United States. Similarly, more than 27 percent of the 100,000 Computer Society members are from outside the United States.

Furthermore, in response to international interest, the Computer Society established a European office in Belgium in 1985 and opened an Asian office in Japan in 1988.

As part of a transnational organization, we have had substantial international participation in our successful standards development work. Consider, for example, IEEE Standard 896.1-1988. Chaired by a British subject with enthusiastic European participation, it has recently been chosen by the US Navy as its future bus standard.

And yet, what happens to our standards after IEEE approval? In accordance with IEEE policy and because IEEE is a standards organization accredited by the American National Standards Institute, the standards are routinely submitted to ANSI for approval.* From there, the path has led (sometimes concurrently) to either the International Standards Organization (ISO) or to the International Electrotechnical Commission (IEC), depending on the nature of the specific standard.

From a larger perspective, looking at this process with a sense of history, several things stand out. In some cases, the process has worked. Some IEEE standards have gone forth to the ISO/IEC arena and done very well. In other cases, the process has not worked.

Now, to provide a balanced appraisal, after more than 10 years of problems, ISO and IEC have established a joint technical group that handles matters of common interest, and ANSI has cleared out the thickets that were blocking access to the international arena.** So now, we have a promise that things are going to get better and, even further, we have some very dedicated people working hard to make it all come true.

What direction to pursue? In view of the above, why not simply continue to press forward on the ANSI/ISO/IEC path? Why should anything else even be considered? There are several reasons to do so. First, all that we have now are promises, not results. Promises have been made before — and broken. A prudent person looking at what has previously taken place might not want to entrust all the IEEE standards to this process.

Second, the ISO/IEC process is a long, slow one. Five years is not an unusual time for documents to stay with those groups. Many documents take even longer. By the time a standard has been through the IEEE process of four years, plus another six months at ANSI, and another four years at ISO/IEC, the original standard is out of date. When this point is raised, there are usually two responses:

(1) The ISO/IEC process should be initiated concurrently with the IEEE process. The danger here is that the IEEE consensus process becomes driven/overridden by the ISO/IEC process and what results may not be satisfactory to the originators.

(2) The ISO/IEC have a "fast-track" process that permits national standards that have had substantial international participation to short cut the five-year process and go directly to ballot. This sounds eminently reasonable and, as such, was tried with a group of IEEE standards that met those criteria. But, the balloting failed to endorse those standards. Reviewing the negative reasons

---

* Section 8.1 of the *IEEE Policy and Procedures Manual* states in part that "... The IEEE shall ... routinely submit all IEEE standards to ANSI for adoption as American National Standards ..."

** For details on ISO, IEC, and the joint technical group, see *Information Technology Standardization, Theory, Process, and Organization* by C. Cargill, Digital Press, 12 Crosby Dr., Bedford, MA 01730.

leaves the reader with something less than a deep appreciation of the technical astuteness of the negative balloters.

Third, the IEEE has a large category of standards that ISO/IEC do not even appear to address. As part of IEEE standards, we produce *Recommended Practices and Guides*. These types of documents are not issued by ISO/IEC except perhaps as technical reports.

Throughout all of this, we should recognize the diversity of interests involved. ISO/IEC have one set of interests, ANSI has another set, and each of the other national organizations has its own. These other interests do not always coincide with the IEEE's. The IEEE is best served when it pursues its own enlightened self-interest and does not act as (or appear to be) a captive of one or more other standards organizations.

Given all that, what is the IEEE interest? Basically, we want to provide consensus-derived standards in a timely manner that will help professionals deal with the problems of the computer industry. As noted above, ISO/IEC does not always fill that need.

Taking that point one step further, we are interested not only in providing the standards that originated with the IEEE, but also in adopting those standards produced by other volunteer standards organizations outside the United States. Certain standards prepared by other organizations have direct value to the IEEE audience and yet, at this time, the only method of acquiring other international standards is via ISO/IEC. And, those other standards organizations have exactly the same problems in processing standards through ISO/IEC as does the IEEE.

What action to take? Given all the above, what should we do? First, we don't want to undo what has already been achieved, so current efforts with ISO/IEC should continue.

Second, recognizing that the IEEE is broad based, we should consider entering into a special relationship directly with ISO/IEC, as has the European Computer Manufacturers Association.

Third, we need to establish a policy at the IEEE level that we will adopt standards from non-US organizations that coincide with the scope of IEEE interests and that have been developed in the same or an equivalent consensus manner as IEEE standards.

Fourth, in accordance with Section 8.1.2 of the *IEEE Policy and Procedures Manual*, we need to establish the criteria for adopting standards and the procedures to implement those criteria.

Fifth, we need to initiate coordination with non-US standards bodies to implement these actions.

Other items should also be considered; for instance, identifying a system for classifying some standards as transnational because of a far-broader-than-US consensus within the IEEE.

The IEEE Standards Board is currently examining this area in terms of policy, criteria, and procedures. Your comments and suggestions on possible solutions or new ideas on these issues would be welcome. For the sake of efficiency, please direct your responses to Andrew Salem, Staff Director, IEEE Standards Department, 445 Hoes Lane, Piscataway, NJ 08855-1331.

PROGRAM COMMITTEE

Fred Lochovsky-Chair	U. Toronto, Canada
Stephen Ades	PA Technology, U.K.
Agustin A. Araya	FMC Corporation, USA
Walter Bender	MIT Media Lab, USA
Elisa Bertino	IEI-CNR, Italy
Andrew Clement	York U., Canada
Panos Constantopoulos	U. Crete, Greece
Christsos Faloutsos	U. Maryland, USA
Steve Feiner	Columbia U., USA
Les Gasser	USC, USA
Sidney E. Harris	Claremont, USA
Mathias Jarke	U. Passau, W. Germany
Robin Jeffries	Hewlett-Packard Labs, USA
Won Kim	MCC, USA
Winfried Lamersdorf	IBM Heidelberg, W. Germany
Larry Lefkowitz	U. Massachusetts, USA
Marilyn Mantei	U. Toronto, Canada
Najah Naffah	Bull S.A., France
Oscar Nierstrasz	U. Geneva, Switzerland
Robert W. Root	Bellcore, USA
Walt Scacchi	USC, USA
Frank Tompa	U. Waterloo, Canada
C.J. van Rijsbergen	U. Glasgow, U.K.
Toyohide Watanabe	Nagoya U., Japan
Carson C. Woo	UBC, Canada
Polle Zellweger	Xerox PARC, USA

CONFERENCE COMMITTEE

Bob Allen-General Chair	Bellcore, USA
Skip Ellis	MCC, USA
Carl Hewitt	MIT-AI Lab, USA
Fred Lochovsky	U. Toronto, Canada
Vince Lum	NPS, USA
Tom Malone	MIT-Sloan, USA
Margi Olson	NYU, USA
Barbara Pernici	Politecnico di Milano, Italy
Dennis Tsichritzis	U.Geneva, Switzerland
Alex Verrijn-Stuart	U. Leiden, The Netherlands

CALL FOR PAPERS

COIS90 - Conference on Office Information Systems

Sponsored by • ACM-SIGOIS • IEEE Computer Society TC-OA

In Cooperation with IFIP 8.4

MIT - Cambridge, MA - April 25-27, 1990

The coordination and support of cooperating computer users is a multi-faceted challenge. Indeed, the issues cover many of the most active areas of computer science. For instance, organizations of intelligent agents may be designed to interact in ways that complement human organizations. New information technologies, intelligent interfaces, AI/knowledge tools, and data models may be combined to provide a wide range of applications such as hypertext, planning systems, multimedia services, and electronic publishing. In addition, these technologies impact individuals, groups, and organizations, and their adoption may depend on user and social design principles. This fifth COIS meeting continues a tradition of multi-disciplinary research on information systems.

- **Coordination Technology/ Cooperative Work**
- **Hypertext/Multimedia**
- **User Interfaces**
- **Social Impact**
- **Electronic Publishing**
- **Information Retrieval**
- **Object-Oriented Systems/Databases**

Information for Authors:
Submit five copies of a double-spaced paper (5000 words maximum), or equivalent electronic submission to:

Dr. Fred Lochovsky
Computer Systems Research Institute, University of Toronto
Toronto, Ontario M5S 1A4
Canada
phone: (416) 978-7441
email: fred@csri.toronto.edu

Submission deadline: November 3, 1989; Notification of Authors: December 30, 1989

Videos:
Tapes for a video presentation relative to themes of the conference are invited. They should be in 3/4 inch or U-matic VHS format (5-15 minutes).
Submission deadline: March 1, 1990

acm Association for Computing Machinery

IEEE COMPUTER SOCIETY

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

If you're looking for the latest in technology, challenging work, and opportunity—we have it!

The next major advance in secure communications is taking shape now at the Electronic Defense Communications Directorate of GTE Government Systems.

If you're ready to apply your abilities . . . there's no better place than GTE Government Systems. We don't look for ready-made solutions. Instead, we encourage you to create, develop your ideas, and find new ways to apply state-of-the-art technology.

Join the Electronic Defense Communications Directorate of GTE Government Systems, and help take communications to a new plateau. Right now, we're looking for professionals interested in the following areas:

Systems Engineering

- Secure Communications
- Communications Protocols
- Concept Definition
- Requirements Analysis
- Trade-Off Analysts
- System Integration

Software Engineering

- Real-Time Communications Networks
- Communication Protocols
- Computer/Communications Security
- VAX/VMS
- Pascal and Assembly Languages
- Integration and Test

Hardware Engineering

- Digital Signal Processing
- High Frequency
- LPI/AJ
- Spread Spectrum
- Modems

All positions require a BS and/or MS in EE, CS, or equivalent experience. Strong communication skills a must.

Candidates must be U.S. citizens with prior DoD experience or the appropriate background to work in the defense community.

At GTE Government Systems, you'll enjoy high visibility, ongoing challenges, and competitive compensation. Our benefits package includes a flexible spending plan, tax-deferred savings plans, a stock purchase plan, educational assistance, and more.

For confidential consideration, please send your resume and salary requirement to: **GTE GOVERNMENT SYSTEMS CORPORATION, Electronic Defense Communications Directorate,** MC-IEC-08, 100 First Avenue, Waltham, MA 02254. An equal opportunity employer, m/f/h/v. U.S. citizenship is required.

GTE Government Systems

Editor: Guylaine M. Pollock, Sandia National Laboratories, 1424, PO Box 5800, Albuquerque, NM 87185; phone (505) 846-0040; electronic mail: gmpollo@sandia.gov.

TC on Computer Graphics focuses on visualization

Gary Laguna, Lawrence Livermore National Laboratory

The past two years, the IEEE Computer Society Technical Committee on Computer Graphics has primarily focused its attention on issues of visualization of information, both scientific and business.

With ever greater computer power available to scientific and business users, the issue of how to efficiently extract, display, and analyze information assumes critical importance. The Computer Graphics TC is seeking to create platforms for visualization scientists and visualization users to exchange information and understand current needs and emerging technologies.

The TC provided several speakers for the successful ONR/NRL Workshop on Visualization in Scientific Computing in May 1988. It was sponsored by the Office of Naval Research and held at the Naval Research Laboratory. The TC assisted the Computer Society in developing congressional testimony regarding possible legislation to allow copyrighting computer terminal scene displays, independent from the software creating the displays.

Larry Rosenblum of the Naval Research Laboratory served as TC chair from 1986 to 1988, and Stephen Levine of Wang has been TC chair since 1988. During their tenures, the TCCG has initiated a wide range of visualization activities for 1989 and beyond:

(1) This special issue of *Computer*, under the theme of "Visualization in Scientific Computing," is sponsored by the TCCG. Gregory Nielson of Arizona State University, a TCCG Executive Committee member, is the guest editor.

(2) Segments on visualization and computer graphics are being generated by the TCCG for Compusat 89, the Interdisciplinary World of Computing satellite symposium. Arthur Olson of the Research Institute of Scripps Clinic, also a TCCG Executive Committee member, is developing the visualization segment. Laurin Herr of Pacific Interface is developing the computer graphics segment. Rosenblum and Anneliese von Mayrhauser of the Illinois Institute of Technology are co-chairing Compusat 89. The event is expected to be held in November or December. Compusat 88, the successful first Computer Society satellite symposium, was conducted last October, with more than 120 sites throughout North America and some 6,000 persons in attendance.

(3) A conference, Visualization 90, is being developed for the fall of 1990 to focus on the visualization needs of both the scientific and business communities. The conference committee has met twice, and an innovative program is planned. Levine is heading the committee coordinating development of Visualization 90 and will serve as general chair.

Levine, the current TC chair, holds a PhD in computer science from Stanford University, is well known for contributions to the computer graphics field, is past president of the National Computer Graphics Association, and has been both vice chair and director of ACM SIGGraph.

Gary Laguna of Lawrence Livermore National Laboratory is TCCG vice chair. In addition to Rosenblum, Levine, Nielson, Olson, and Laguna, other members of the TCCG Executive Committee are Maxine Brown of the University of Illinois at Chicago, Michael Danchak of the Hartford Graduate Center, John Sibert of George Washington University, and Mark Skall of the National Institute of Standards and Technology.

Membership in the TC is free to any Computer Society member. To obtain additional information or a TCCG application, contact Gary Laguna, PO Box 808 (L-125), Lawrence Livermore National Laboratory, Livermore, CA 94550.

Computer Society Press offers new videotape series

Gene Falken, Director,
Computer Society Press

The Computer Society Press is producing nine new videotapes that will be on sale soon. Materials should be available in September.

Based on state-of-the-art technology recently presented through lectures and visuals at IBM's US Education Center in Thornwood, N.Y., the tapes were made available for CS Press distribution through the generous assistance of IBM.

Ez Nahouraii of IBM, the editor-in-chief of CS Press, developed the program in collaboration with Frederick Petry of Tulane University, the newly appointed CS Press editor for media.

Running approximately 150 minutes, each tape provides a depth of treatment not usually found in typical one-hour presentations. A book of notes and visuals accompanies each tape.

The tapes cover a variety of topics. Topics and presenters include:

- *Neural Networks—Current Applications and Directions,* Bruce Shriver, University of Hawaii and editor-in-chief of *Computer*
- *Parallel Processing—Architectures and Directions,* William Dally, Massachusetts Institute of Technology
- *Improving Software in a Measurable Way,* Victor Basili, University of Maryland
- *Software Risk Management,* Barry Boehm, TRW
- *Communications and Synchronization in Parallel Processing Systems,* Harold Stone, IBM
- *Distributed Software Engineering,* Sol Shatz, University of Illinois
- *SQL-2 Standards: ISO and ANSI,* Philip Shaw, IBM
- *Distributed Supercomputing Networks,* Jordan Becker, IBM

- *Software Reliability Measures: Guiding Software Development for Quality and Cost Effectiveness,* John Musa, AT&T Bell Laboratories

Previews of the videos, screened in mid-May at the International Conference on Software Engineering in Pittsburgh, prompted a great deal of positive reaction and on-the-spot ordering. Other reviewers have seen the entire program and are equally enthusiastic.

The videos are in VHS format. Each is tentatively priced at $89 list and $69 for members and is accompanied by printed notes and visuals. For non-US prices, add 20 percent. If the final prices are higher than those stated, all orders received through September 30 will be honored at the above prices.

All tapes are sold with a full refund privilege if the user is dissatisfied in any way.

Order from Customer Service, Computer Society Press, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720, phone (800) CS-BOOKS (in California, dial (714) 821-8380), fax (714) 821-4010.

The CS Press has other video projects under consideration and would welcome suggestions or advance copies of new tapes of good quality focusing on computer science and engineering.

Rate correction

Beginning in 1990, the member subscription rate of *IEEE Expert* will rise to $18 a year from the present $12, increasing $6 not the $4 stated on p. 81 of the July issue of *Computer*.

The IEEE Computer Society Board of Governors adopted the new rate at its May 17 meeting in Pittsburgh to offset the increased cost of producing 210 additional pages in the magazine in 1990, when it will change from quarterly to bimonthly frequency.

UPDATE

Contributions to Update are welcome. Send news of industrial or university research and of public policy or professional issues to Update Editor, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720, or to Editor-in-Chief Bruce D. Shriver, Vice President for Research, University of Southwestern Louisiana, Drawer 42730, Lafayette, LA 70504.

IEEE-USA workshop applauds DRAM consortium

The recently announced US Memories initiative to jointly manufacture DRAM chips in the US must be supported "if the US is to reestablish a strong technology and manufacturing base." This was among the conclusions reached by the members of a June 27 workshop on the role of a US DRAM initiative in revitalizing the US electronics base, sponsored by the Institute of Electrical and Electronics Engineers-United States Activities.

In addition to endorsing the broad outlines of the USM initiative, the workshop members called on the US government to support the initiative and endorsed establishment and support of similar ventures. Some 70 participants from industry, government, and academia joined in the workshop.

The USM initiative announced June 21 proposes a $1 billion joint venture to manufacture DRAM chips in the United States. USM participants are Advanced Micro Devices, Digital Equipment Corporation, Hewlett-Packard, Intel, LSI Logic, and National Semiconductor. IBM's recently announced 4-megabit chip is expected to be the cornerstone of USM's entry into the world market. The participants are now studying the feasibility of the venture, which is also subject to regulatory approval by the US government.

The workshop group stated that USM achieves "virtual vertical integration," allowing several US firms "to achieve many of the benefits of integration and thus to compete more effectively with major foreign producers of DRAMs that are truly integrated firms."

The group noted that the US government should

- recognize the importance of a "hospitable antitrust environment" to attract private US capital;
- support an adequate administrative capability to monitor trade flows and promptly and effectively implement existing government policies and rules; and
- equalize capital costs and risks for US participants relative to foreign competitors in global markets.

In general, the group stated, "government must support industry-led initiatives for dual use (both civilian and military) technologies by providing an environment and incentives for US firms that are economically attractive and reduce risks to US participants."

According to the statement, USM and similar initiatives can provide significant additional US-owned, secure sources of DRAMs; provide additional markets for the semiconductor manufacturing equipment sector in US industry; and complement the overall efforts and contributions of Sematech, the semiconductor manufacturing industry consortium.

The workshop participants also endorsed "establishment and support of ventures that help ensure an enduring independent, competitive US DRAM manufacturing capability . . . (including) not only fabrication facilities, but also associated research, design and development facilities, as well as tool development and manufacturing; (and) ventures that consist of as few as one US-owned domestic DRAM manufacturer."

"If available funds are significantly constrained, then any `on budget' elements of direct government support for US memory manufacture (should) be initially limited to DRAMs to ensure their adequate funding," the group concluded.

The statement containing the participants' consensus was announced at a June 28 press conference by John M. Richardson, chair of the IEEE-USA Committee on Communications and Information Policy, and Alan K. McAdams, chair of the IEEE-USA Technology Leadership Subcommittee. Richardson also presented the statement to Rep. Doug Walgren (D-Pa.), chair of the House Science, Research, and Technology Subcommittee.

US losing supercomputer lead: IEEE supercomputing group

The United States is close to losing its lead in supercomputers, according to the Scientific Supercomputer Subcommittee of the IEEE-USA.

In a revision of an earlier report, the group noted recent developments that it feels demand government action, including NEC's announcement that it is producing supercomputers faster than those made by Cray Research and the closing of Control Data's supercomputer business, ETA Systems. Subcommittee chair Sydney Fernbach also cited problems with government procurement during the 1970s.

The subcommittee feels the federal government should

- commit itself a well-defined, long-range national program to maintain leadership in scientific supercomputing development;
- encourage both the continued development of scientific supercomputers and their availability by adopting such measures as direct support for the development of new supercomputers and easing of antitrust rules for research and development consortia;
- support technologies needed for the development of supercomputer systems;
- designate a lead organization to coordinate the federal agencies that depend on supercomputers; and
- establish and extend national performance goals for the development of scientific supercomputer systems.

Fernbach said that the recent US supercomputer trade agreement with Japan could work to the disadvantage of the United States as Japanese supercomputers become more competitive with US machines. He also criticized the lack of industry and government cooperation in supercomputers.

The subcommittee asserted that supercomputers are essential for superior performance and lower costs for a wide range of high and medium technology products. It also tied the US lead in supercomputer technology to US military and economic security and international competitiveness.

Copies of the revised report are available from IEEE United States Activities, 1111 19th St. NW, Suite 608, Washington, DC 20036-3690, phone (202) 785-0017, attn. Heidi F. James.

News briefs

Computing education groups merge. The International Council for Computers in Education and the International Association for Computing in Education have merged to form the International Society for Technology in Education. The merged organization and its geographic affiliates will include more than 40,000 members.

ISTE reportedly seeks to improve the quality of education through technology by providing support for computer-using educators, facilitating international sharing of information and resources, and encouraging research, innovation, and organizations.

ISTE will publish *The Computing Teacher* journal, *Update* newsletter, *The Journal of Research on Computing in Education*, and several other periodicals and monographs.

ISTE will maintain membership in the American Federation of Information Processing Societies and the Institute for Certification of Computer Professionals. It has also applied for category C membership in UNESCO.

ISTE can be reached at 1787 Agate St., Eugene, OR 97403, phone (503) 686-4414.

NSF report warns of education crisis. A report issued after a series of National Science Foundation-sponsored workshops warns of an impending crisis due to an alarming lack of attention to undergraduate education in the sciences, mathematics, and engineering.

The *Report on the National Science Foundation Disciplinary Workshops on Undergraduate Education* cites a preoccupation by faculty with research rather than teaching. Workshop participants stressed the need for comparable support for innovative teaching.

One solution offered was the establishment of grants for outstanding teaching and curriculum, which the report recommended.

The report, publication No. NSF 89-03, is available from the Division of Undergraduate Science, Engineering, and Mathematics Education, Room 639, 1800 G. Street NW, Washington, DC 20550.

Cornell takes on high school supercomputing program. The Cornell National Supercomputing Facility has taken over the national SuperQuest competition from ETA Systems. The competition introduces high school students to supercomputing. Winning teams will conduct supercomputer studies of airplane wings, eye surgery, traffic lights, black holes, robots, medical X-rays, "brainlike" computers, colliding galaxies, advanced maps, concert halls, and island ecology.

AEA, SIA to press reform of dumping rules

Senior executives of the American Electronics Association and the Semiconductor Industry Association have agreed to press for fundamental reforms in the international dumping code and the US antidumping law. The joint committee proposed rewriting principles and practices to deter dumping and submitting that proposal to the US government for inclusion in US law and administrative practices.

The group also agreed to urge the US government to present the proposal for inclusion in the General Agreement on Tariffs and Trade.

The committee drafted a policy statement on market access in Japan by foreign semiconductor manufacturers. It recognized increases by several major Japanese users, but noted the lack of improvement in overall market share.

Despite ongoing negotiations between US producers and Japanese users, the group advised retaining sanctions until the administration has reviewed their overall effectiveness.

The AEA/SIA formed the joint steering committee in the fall of 1988 to address international and competitive issues, emphasizing implementation of the US/Japan Semiconductor Trade Agreement of 1986.

PRODUCT REVIEWS

Editor: Richard Eckhouse, MOCO, Inc., PO Box A, 91 Surfside Rd., Scituate, MA 02055; Compmail+, r.eckhouse

TeXnology on the IBM PC

T. L. (Frank) Pappas, Intermetrics

The TeX typesetting system, developed at Stanford under the direction of Donald Knuth, is described in *The TeXbook*. In the introduction, Knuth writes that TeX (which he says rhymes with blecchhh) was originally developed "... for the creation of beautiful books—and especially for books that contain a lot of mathematics." While TeX has been very successful for these purposes, it's being used successfully in other areas as well.

TeX has been around for more than ten years, but many people have never heard of it. It's available on a wide range of computers including CDC, Cray, IBM (mainframe and PC), Atari, Macintosh, DEC, and Sun. In this review I'll look at the new IBM PC implementations. I'll also look at other products—some TeX-related, some not—that will make your TeXing more effective.

A different look

As a clear demonstration of what a non-expert can do with TeX, look at this review. The editors of *Computer* let me produce it in TeX and suspended the magazine's rules of style so I could show you how TeX handles a variable baseline, font changes, and so on. I printed the review on a Brother HL-8e laser printer and submitted a copy to *Computer*. What you see is a photoreproduction.

TeX output produced on a typesetter looks much better than what you see here. Of course, most of you don't have easy access to a typesetter, so you'll be using a laser printer. Combined with the fact that there's much more to TeX than I could possibly cover in this review, this approach lets you see what you can produce using TeX, without my having to discuss all of the features that I use.

If you do need to produce typeset quality, organizations such as the American Mathematical Society will provide this service for you on a per-page cost. The TeX Users Group (see the sidebar) publishes a newsletter, *TUGBOAT*, which has ads from organizations providing this service.

The only portion of this column not produced by TeX is the department heading and the footer. These can be produced by TeX, but it was more convenient to use the standard department heading and to give the editors the freedom to place and number the pages.

TeXnology overview

I'll begin with an overview for readers not familiar with TeX, but first let me warn you that it takes a while to become proficient. Although I've been using it on and off for more than a year, I'm not an expert, but I can use it effectively. For me, it's been worth the effort spent learning it.

I'll share some of my TeX experiences with you to give you an idea if it's for you. For example, I'm using TeX to write a book on Ada software reuse. The book contains many Ada code fragments, which I want to make as presentable as possible. As I'll describe later, TeX is simplifying my preparation of the book and will make it very easy to deliver a finished product to my publisher. I've also used it to prepare view graphs and white papers.

A companion program, METAFONT, creates fonts for TeX. Except for a few graphics described later, all of the text and symbols used in this review are in fonts created by METAFONT. Most of you won't need it since you'll get the needed fonts when you purchase TeX.

TeX, METAFONT, and some related utilities are written in a structured documentation language called Web, which combines elements of TeX and Pascal. One Web processor, Weave, generates a TeX input file that you use to generate a typeset version of a Web program. Another Web processor, Tangle, generates a Pascal file which, when compiled, is the executable program. This allows the TeX documentation and Pascal source for TeX, METAFONT, etc. to each be combined in a very readable program. Knuth calls this style of programming *literate programming*.

The Web source for TeX, METAFONT, etc., a Pascal version of Tangle and the TeX source for *The*

TeXbook and *The METAFONTbook* are available for a modest cost. Versions of the programs are also available in C. I'll tell you later where you can get them.

TeX input. Since it's not a WYSIWYG (what you see is what you get) system, you create an ASCII input file for TeX. The file contains both the text you want typeset and embedded commands to direct TeX.

A command consists of the command trigger \ followed by a word or a special character. For example, to change to typewriter text or to italics, you use the `\tt` and `\it` font change commands, respectively. To typeset '`this typewriter text` *is followed by italicized text* and then by roman text', you simply enter '`{\tt this typewriter text \it is followed by italicized text\/} and then by roman text`'.

The braces limit the scope of the font changes. They serve a purpose similar to begin-end brackets in Algol-like languages. The `\/` command provides an italic correction, that is, it adds a little bit of extra white space when switching from a slanted to an unslanted font.

TeX does a nice job of dealing with accents. For example, to produce 'George Pólya and Gabor Szegö', you type '`George P\'olya and Gabor Szeg\"o`'. Dotless i and j are provided for with the commands `\i` and `\j`, respectively. These let you typeset words like 'mı̄nŭs' without the dot over the i getting in the way as it does in 'mi̇̄nŭs'.

TeX does lots of little things for you that you don't even have to think about. For example, ligatures are handled without any action on your part. Typing `flavor` gives you flavor instead of flavor.

Justification and hyphenation. As the next few paragraphs illustrate, TeX does a nice job of justifying text. TeX gets really fussy when it starts justifying. If it can't do a good enough job of line breaking, it inserts a little black box at the end of an offending line. To remove the box, you might have to change a word or two, insert a discretionary hyphen, or even rewrite the paragraph.

TeX hyphenates words, when necessary, to create the best line breaks possible. Because of its conservative approach to hyphenation, TeX rarely hyphenates a word incorrectly, even though its hyphenation algorithm is not foolproof. TeX won't hyphenate words that have different hyphenations based on use. For example record can be hyphenated as *re·cord* or *rec·ord*, depending on whether it is used as a verb or a noun.

TeX Users Group

The TeX Users Group (TUG) publishes *TUGBOAT*, a magazine that you get as a member. In addition to containing all the latest news in the TeX world, articles include macros and examples that you can use to increase your effectiveness with TeX. The macros are usually available electronically. The modest membership fee is worth it.

You should also ask TUG about *TeXhax*, a free weekly electronic magazine containing useful macros and ideas. You can send TeX questions to *TeXhax* and other readers will usually answer.

You can request information about TUG through regular mail by writing to

TeX Users Group Office
PO Box 9506
Providence, RI 02940-9506

by electronic mail at their Internet account

TUG@Math.AMS.com

or by calling them at

(401) 751-7760

That's all right, because TeX doesn't need to hyphenate too often since it typesets a paragraph at a time, not a line at a time. The end of a line is treated like a space, which causes TeX to view the paragraph as one long line. Therefore, the last line in a paragraph can affect how the first line looks.

When justifying, TeX attempts to break a paragraph into lines by adjusting the space between words so that no line contains too much white space. When a ragged right edge is used, TeX attempts to prevent any line in the paragraph from having too much white space at the end by adjusting the white space at the end of the line. In either case, TeX inserts a hyphen only when adjusting the white space alone doesn't work.

Sometimes TeX can't do a good enough job of keeping the white space under control because it doesn't know how to hyphenate a particular word. You can use a discretionary hyphen, as in `rec\-ord`, to tell TeX how to hyphenate a particular instance of `record`. If you use record a lot, and only as a noun, you can write `\hyphenation{rec-ord}` to tell TeX how to hyphenate all occurrences of `record`.

Fonts. TeX usually comes with a large collection of fonts called the Computer Modern (CM) fonts. These include 75 different fonts ranging in size and style, with up to six different magnifications. This review is written using CM 9-point roman for normal text. The title of the review uses CM 18-point roman.

As this strange-looking paragraph shows, it's easy to switch fonts and sizes within a paragraph. It starts in 9-point roman, then switches to 11-point san serif, then 10-POINT SMALL CAPITALS, then *10-point italic san serif*, and then 10-point san serif. The range of sizes allows you to write text that is smaller and smaller and smaller and smaller.

If you're not satisfied with just the standard fonts, you can create your own using METAFONT. You can also use it to create logos and geometric designs.

Math typesetting. Typesetting mathematics by hand is considered difficult, but it's considerably easier using TeX. Essentially, TeX has two types of math mode—text style and display style. Text style is used for generating formulas embedded in text, such as $\sum_{n=1}^{m} x_n$. Display style is used for displaying stand-alone formulas, such as

$$\sum_{n=1}^{m} x_n$$

Notice that, by default, TeX places the limits of the sum at the top and bottom of the summation symbol when the formula is displayed, but puts them to the side when the formula is embedded. Again by default, TeX uses an italic font suitable for typesetting formulas and suppresses spaces, since they don't normally appear in formulas.

The text style version of this formula is written as

```
$\sum_{n=1}^m x_n$
```

where $ is used to enter and leave math mode. The `\sum` command gives you the summation symbol. The character _ gives you a subscript, while ^ gives you a superscript. The display style version uses $$ instead of $.

The figure on the right contains other examples of formulas typeset in TeX. These might look complicated to typeset, but they are actually pretty simple once you learn more about TeX. In fact, if you write formulas like these, you can easily typeset them in TeX!

$$\begin{aligned}\left(\int_{-\infty}^{\infty} e^{-x^2}\,dx\right)^2 &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-(x^2+y^2)}\,dx\,dy\\ &= \int_0^{2\pi}\int_0^{\infty} e^{-r^2}\,r\,dr\,d\theta\\ &= \int_0^{2\pi}\left(-\frac{e^{-r^2}}{2}\bigg|_{r=0}^{r=\infty}\right)d\theta\\ &= \pi\end{aligned}$$

$$\begin{tikzcd}
 & & & 0 \arrow[d] & \\
0 \arrow[r] & \mathcal{O}_C \arrow[r, "\iota"] \arrow[d, equal] & \mathcal{E} \arrow[r, "\rho"] \arrow[d, "\phi"] & \mathcal{L} \arrow[r] \arrow[d, "\psi"] & 0 \\
0 \arrow[r] & \mathcal{O}_C \arrow[r] & \pi_*\mathcal{O}_D \arrow[r, "\delta"] & R^1f_*\mathcal{O}_V(-D) \arrow[r] \arrow[d, "\theta_i\otimes\gamma^{-1}"] & 0 \\
 & & & R^1f_*\big(\mathcal{O}_V(-iM)\big)\otimes\gamma^{-1} \arrow[d] & \\
 & & & 0 &
\end{tikzcd}$$

Examples of formulas typeset in TeX.

Macros. The power and simplicity of TeX can be extended through its macro capability. For example, if you were writing a paper and had to mention Pólya frequently, it would be tiring and error prone to type `P\'olya` each time. To get around this you can write the macro

```
\def\Polya{P\'olya}
```

and then write `\Polya` to get Pólya.

TeX macros can have parameters, so you can write the following macro to accept the limits, index, and sequence of a sum:

```
\def\sums#1#2#3#4{
      $\sum_{#1#2}^#3#4_#1$}
```

Then you can write `{\sums n 1 m x}` to get $\sum_{n=1}^m x_n$, the sum we had earlier, or `{\sums i a b f(x)}` to get $\sum_{i=a}^b f_i(\mathrm{x})$, an entirely different sum.

The previous sentence illustrates how TeX adds a little bit of interline spacing to accommodate the ascender of the letter 'b' in the upper summation limit. It also illustrates why you should avoid embedded formulas.

The macro language is powerful enough to allow registers (variables), conditional tests, looping, and recursion. The *The TeXbook* contains many examples of macro programming.

You don't really use pure TeX, since it's more a collection of building blocks. Instead you use a macro package, so that commands are macros. TeX usually comes with three such packages—Plain TeX, LaTeX, and AMS-TeX. The latter packages, like others, redefine some Plain TeX macros and add new ones to supplement those that haven't been altered.

Unless you're writing papers for the American Mathematical Society, you won't be interested in AMS-TeX, so I won't mention it further. While many people use Plain TeX, I think most readers will be interested in LaTeX or *TeXT1*, a separately purchased product, both of which I'll discuss in later sections.

DVI files. When you run TeX on a file, you don't get output for your printer, instead you get a device-independent file. A DVI file needs to be processed by a driver program, which can be either a previewer for viewing a typeset document at a terminal or a printer driver.

Actually, DVI files are portable unless you use the TeX `\special` command (described below). You can create a DVI file on a PC, send it to, say, a VAX/VMS system and print it there, and vice versa.

For my book, I'm going to produce a DVI file and ship it off to my publisher, who will use a driver specific to whatever typesetter he decides to use. This saves my publisher the expense of typesetting the book from scratch and saves me the effort and aggravation of proofreading what a typist entered.

LaTeX macro package. This popular macro package provides a very friendly TeX interface, especially to the TeX novice. For those of you familiar with the Scribe document preparation system, Leslie Lamport patterned LaTeX after it.

LaTeX comes with predefined styles for writing articles, reports, books, and letters. A special version of LaTeX lets you prepare slides.

```
\begin{itemize}
\item In an itemized list ... with a bullet.
\item You can have enumerated lists as well.
\begin{enumerate}
\item The items are numbered for you.
\item You can add ... that follow.
\begin{enumerate}
\item Enumerated lists can be nested.
\item Distinct counters are are maintained.
\end{enumerate}
\end{enumerate}
\item These lists are easy to use.
\end{itemize}
```

- In an itemized list, each item is marked with a label. By default the first list level in an itemized list is marked with a bullet.
- You can have enumerated lists as well.
  1. The items are numbered for you.
  2. You can add items in the middle without having to renumber the ones that follow.
     (a) Enumerated lists can be nested.
     (b) Distinct counters are maintained.
- These lists are easy to use.

```
\begin{CO}\label{Anticipated}
Do not use the predefined exceptions to
handle anticipated conditions.
\end{CO}

\begin{Discussion} The reusability of the
following subprogram is compromised
because of its inappropriate dependence on
{\tt Constraint\_Error}:
\begin{Example}
\Procedure Increment
            (Counter : \In \Out Integer) \is
\Begin
  Counter := Counter + 1;
\Exception
  \When Constraint\_Error =>
    Counter := Integer'First;
\End Increment;
\end{Example}
The potential problem exists because the ...
\end{Discussion}
```

Guideline 11.10 (Composability) *Do not use the predefined exceptions to handle anticipated conditions.*

Discussion: The reusability of the following subprogram is compromised because of an inappropriate dependence on `Constraint_Error`:

```
procedure Increment
            (Counter :  in out Integer)  is
begin
  Counter := Counter + 1;
exception
  when Constraint_Error =>
    Counter := Integer'First;
end Increment;
```

The potential problem exists because the subprogram is not written for reuse in a hostile environment. Its execution may be erroneous if it is compiled without runtime checking—maybe it appears in a ...

LaTeX input (left) and LaTeX result (right).

With LaTeX, you can easily divide papers, books, and reports with sectioning commands. For example, `\section{Exception Handlers}` will produce a numbered section title with a slightly larger font and reasonable spacing before and after the title. If it's the second section in chapter 11, it will be numbered 11.2, so you can easily move sections around. Chapter, subsection, and paragraph titles are also appropriately displayed and numbered. If you request them, a table of contents, figures, and tables can each be generated directly from the sectioning commands. There are even facilities for creating an index or a bibliography.

LaTeX has many useful environments for displaying text, drawing pictures (curves, circles, ellipses, lines, arrows), making tables, floating figures, etc. You can also define your own environment.

I'll give two examples of how easy it is to use LaTeX. The first example in the figure at the top of the page shows how easy LaTeX makes it to create itemized and enumerated lists. You don't have to indent the list items or number them. LaTeX does that for you. And you don't have to deal with hanging indents and tabs the way you do in a WYSIWYG system.

The second example illustrates several features, which I used for my book. Basically, the book consists of a collection of four kinds of guidelines, one of which is **Composability**. I made a minor change to a LaTeX file and then used two predefined LaTeX environments to create a new CO environment as follows:

```
\newtheorem{Guideline}%
{\gf Guideline\/}[chapter]
\newenvironment{CO}%
{\begin{Guideline}%
[Composability]\sl}%
{\end{Guideline}}
```

It generates chapter-relative guideline numbers, switches to a slanted font, emphasizes the heading, and boxes the guideline.

I defined the Example environment for displaying Ada code. This environment switches to a monospacing typewriter font and centers the example. The Discussion environment, which I also defined, slightly indents the text of the discussion and adds ***Discussion:*** to the front of it.

For each Ada reserved word I defined a macro using its capitalized name to simplify the text entry of the examples. For example, the macro for procedure is `\def\Procedure{\ts procedure}` where `\ts` is simply a macro defining the font I'm using.

One last feature in the example is the use of the `\label` command. I can generate a reference to the guideline or the page on which it appears using the label name and the `\ref` or `\pageref` command. I can write 'See Guideline `\ref{Anticipated}` for …' and get 'See Guideline 11.10 for …'.

Learning TeX. One of the reasons it takes a while to learn TeX is that *The TeXbook* is not an introductory manual. Knuth tried to make it both an introduction and reference manual by using the dangerous-bend symbol to mark paragraphs that you skip on a first reading. But some of what you want to use TeX for, even as a beginner, is in these paragraphs. Unfortunately, they often discuss one topic in terms of topics appearing in later paragraphs, just like a reference manual. Once you learn TeX, that's no problem because of Knuth's writing ability, but it's a rough way to learn TeX.

Because it takes a while to learn TeX, I recommend you start with LaTeX or *TeXT1*, instead. Eventually, you will want to learn Plain TeX to get the most out of TeX. An inexpensive 90-page collection of TeX notes, titled *TeX Made Easy*, will help. Prepared by Zar Limited, it's patterned after a five-day course on Plain TeX. At $15 it's a must.

Special effects. To illustrate the kinds of special effects you can create using TeX, I've gone to a single-column format and turned on justification. Notice that the columns from the three-column format are balanced before and after the single-column format. The `\onecolumn` command switches to single column, while the `\threecolumns` command switches back to three columns. There's also a `\twocolumns` command.

The following paragraph comes directly from *The TeXbook*. The dangerous-bend symbol that starts the paragraph signals a paragraph that should be skipped on a first reading. It's part of a font created for *The TeXbook* using METAFONT. The major special effect, of course, is the embedded circle. The paragraph gives an overview of how it is created.

☡ It's possible to control the length of lines in a much more general way, if simple changes to `\leftskip` and `\rightskip` aren't flexible enough for your purposes. For example, a semicircular hole has been cut out of the present paragraph, in order to make room for a circular illustration that contains some of Galileo's immortal words about circles; all of the line breaks in this paragraph and in the circular quotation were found by TeX's line-breaking algorithm. You can specify an essentially arbitrary paragraph shape by saying `\parshape=<number>`, where the `<number>` is a positive integer n, followed by $2n$ `<dimen>` specifications. In general, '`\parshape=`n i_1 l_1 i_2 l_2 $\ldots$ i_n l_n' specifies a paragraph whose first n lines will have lengths l_1, l_2, …, l_n, respectively, and they will be indented from the left margin by the respective amounts i_1, i_2, …, i_n. If the paragraph has fewer than n lines, the additional specifications will be ignored; if it has more than n lines, the specifications for line n will be repeated ad infinitum. You can cancel the effect of a previously specified `\parshape` by saying '`\parshape=0`'.

The area of a circle is a mean proportional between any two regular and similar polygons of which one circumscribes it and the other is isoperimetric with it. In addition, the area of the circle is less than that of any circumscribed polygon and greater than that of any isoperimetric polygon. And further, of these circumscribed polygons, the one that has the greater number of sides has a smaller area than the one that has a lesser number; but, on the other hand, the isoperimetric polygon that has the greater number of sides is the larger.
[Galileo, 1638]

TeX shortcomings. I've only noticed one real shortcoming with TeX, and that's its lack of graphics support. The graphics that appear later in this review first had to be converted to PCL to be printed on the laser printer. To include the PCL file 'running.pcl' in the review, I had to insert the line `\special{hp:plotfile running.pcl}`. Unfortunately, `\special` varies slightly from implementation to implementation.

I found one minor shortcoming. TeX can't determine the size of a graphic image you import with `\special`, so you have to know how big it is and use TeX commands to reserve space for it. You can place the graphic in a figure and float the figure if you like, but the ease of this depends on the macro package you're using. Floating a figure is easy in LaTeX and *TeXT1*, except that in multicolumn format, *TeXT1* assumes the figure is as wide as the page.

TeX products

Now that you've had a glimpse of what TeX's all about, it's time to look at the PC implementations and related products. Since I've already covered the basics of TeX, these reviews will be somewhat small, just giving the strengths and weaknesses of each product.

There are six TeX implementations for the PC, but only four of them were ready for review. You can get most TeX products directly from TUG. With the exception of *TeXT1*, TeX vendors also offer the TeX-related products and the books. The prices I've quoted are directly from the vendors. Each vendor offers substantial discounts for package deals. A company whose name is also its telephone number, **1-800-USA-BOOKS**, claims you can save money by buying through them.

For this review I used a 386 computer with 640 Kbytes of conventional memory, several megabytes of extended memory, a large capacity (330-Mbyte) hard drive, and a multiscan monitor. You can get by with a PC XT and about 5-8 Mbytes of disk storage (depending on the implementation and options you choose), but I strongly recommend a 286 or 386. You need graphics if you want to use a previewer.

While you can get a dot matrix printer driver, they are painfully slow. Unless you're going to use at least a laser printer, forget using TEX.

To avoid some redundancy in the reviews, I need to say a few words about drivers. A printer driver should honor the `\special` command. It should provide options to print ranges of pages and to print in reverse order. It should also let you substitute one font for another. The printer drivers reviewed below, which are for the HP LaserJet Plus and LaserJet II, have these options.

Similarly, a previewer should let you specify a starting page, magnify or shrink an image, and easily move around on a page. It should also let you substitute one font for another. The previewers reviewed below have these options. While they all provide Hercules graphics, EGA, and VGA support, none of the previewers support higher resolutions.

Each TEX implementation comes with a collection of TFM files. TEX uses these device-independent font files to create DVI files. PK files, which contain font information for printing at a particular resolution, come with your driver. A printer driver and previewer can share a set of PK files. The PK files take up several megabytes, but you can save some space by removing the files for fonts, sizes, and magnifications you don't intend to use.

Reader service information has been collected and placed on the last page of this review. Each company has one reader service number.

PCTEX. This TEX implementation, by Personal TEX Inc. (PTI), is currently the best. It is at least as fast as any of the others and it allows you to use extended memory through a RAM disk. It comes with TEX, LATEX, and AMS-TEX.

The installation documentation is clear. The elementary introduction to TEX included should help you take your first TEX steps. It sells for $249 and requires 2.3 Mbytes of disk space. I used it for this review.

PCTEX also has a shell for invoking TEX and your editor, previewer, and printer driver. There's no problem with the shell if you only have a C drive, but when you have more than one, the shell doesn't work if it's not on the same drive as your input files. This problem has existed for at least two releases, so it probably won't be fixed soon.

PTILaser/HP is a printer driver for the LaserJet. It sells for $195 and requires 327 Kbytes of disk space. The CM files require 2.3 Mbytes of disk space. All fonts and sizes are not available at all magnifications, but the ones you need are probably there. However, PTILaser/HP didn't work with any of the compatible printers that I tried.

PTIView is a previewer. It sells for $149 and requires 222 Kbytes of disk space. It comes with the same collection of CM files as PTILaser/HP. For the most part, the commands are well thought out. For example, the arrow keys let you move around a page, the return key or plus key moves one page forward, and the minus key moves one page backward.

pcMF is an implementation of METAFONT. It sells for $195 and requires 1.6 Mbytes of disk space. While it works without any problems, its documentation is somewhat terse. Since *The METAFONTbook* suffers from the same problems as *The TEXbook*, the poor documentation doesn't help matters. If you want to use METAFONT, be prepared to spend a considerable amount of time learning it.

PTI also sells the Web source for TEX, METAFONT, and the utilities. It comes on 10 diskettes in compressed form and sells for $25. The documentation source comes on six diskettes and sells for $25. You might consider getting the documentation source to see how Knuth used TEX to typeset his books.

μTEX. This version of TEX, by ArborText, would be my recommendation if it supported extended memory. Its documentation could be better—you just get an installation manual. Other than that, I have no complaints about it. According to ArborText, μTEXwill support extended memory by Fall. μTEX costs $249 and requires 4 Mbytes of disk space. AMS-TEX and a slide-making version of LATEX are available together as a separate product.

μTEX has a preliminary previewer built into it. As TEX processes your file, you can use the previewer on the pages already created. Unfortunately, flaws make it unusable. The previewer uses GF files instead of PK files. (METAFONT actually creates TFM and GF files. A utility file creates PK files from GF files.) It doesn't have enough GF files to support the fonts, sizes, and magnifications you might want. When you use a font that it doesn't support, you get empty spots. For example, the chapter titles in my book can't be displayed, even though they are in the default fonts and sizes used by LATEX. If this previewer eventually uses PK files, it will make a nice addition to μTEX. For now, if you don't use the previewer, you can save about 1.1 Mbytes of the 4 Mbytes by getting rid of the GF files.

Included with μTEX is a markup language called Mappit, patterned after SGML. Mappit consists of an executable file and a collection of macro files. It comes with a considerable amount of well-written documentation. I didn't use it, so I don't know how well it works, but it looks interesting. It takes up 606 Kbytes of disk space.

ArborText makes the best drivers I've found. For example, DVILaser is the only one of the printer drivers that works with the LaserJet compatibles I used. It also allows you to put a date and time stamp with some optional text in the page header. It costs $225 and requires 1 Mbyte of disk space, with 2.3 Mbytes more for the CM files. I used this driver to print what you are reading.

DVILaser's utilities and examples account for the large disk space re-

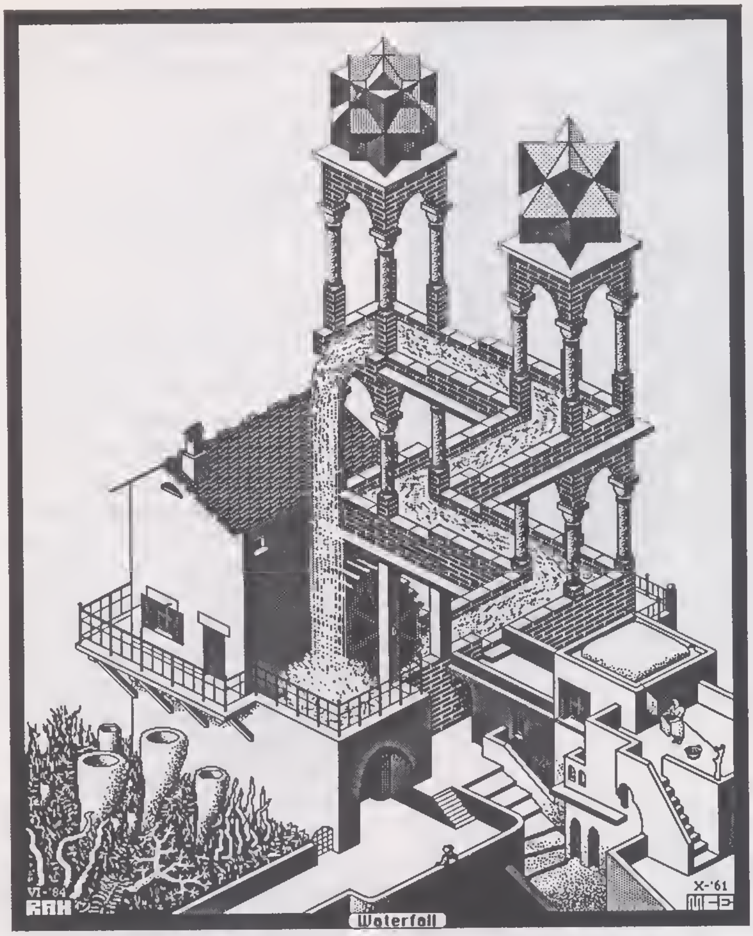

MacPaint to PCL conversion. (Digitized image supplied by ArborText, based on M. C. Escher's "Waterfall" ©1989 M. C. Escher Heirs/Cordon Art, Baarn, Holland.)

quirement. In particular, it comes with Painthp, a program that converts MacPaint files to PCL files. The PCL version of the well-known figure at the top of this page was created using Painthp.

ArborText's Preview is an outstanding previewer, even though its command structure could be more natural. It has several features that the other previewers don't. A scroll mode makes it easier to move through your DVI file, and a two-page view lets you see pages side-by-side. It also has a forward and backward string scan capability. A feature that saved me a lot of time preparing this review is the horizontal and vertical pica rulers. They helped me determine how much space to leave for figures. Preview costs $149 and requires 1 Mbyte of disk space.

Two sets of CM files come with Preview. One is the same set that comes with DVILaser. Preview uses the second set to give the best resolution of any of the previewers. If you don't need this better resolution, you can save about 1 Mbyte of disk space by removing this set and modifying Preview's font configuration file.

CTeX. This TeX version, by Micro Publishing Systems, costs $189 and requires 2.19 Mbytes of disk space. It comes with AMS-TeX and the slide-making version of LaTeX.

Micro Publishing's TeXPrint/HP LaserJet driver didn't work with my compatibles. That's a shame, because it has a feature I really like. You can select even-page or odd-page printing. This way you can create a double-sided copy by printing the odd pages first, then placing the printed pages back in the input bin and selecting even pages with reverse order printing.

This driver costs $189 and requires 260 Kbytes of disk space. It comes with the most extensive set of CM files, with all magnifications supplied. This requires 4.41 Mbytes of disk space, but you don't have to install all of it, since at each of the magnifications you are given the option to skip installing it.

Micro Publishing also sells TeX-Write, an integrated editor and shell, for $149. It requires 130 Kbytes of disk space. It has a few nice features, but unless you're using Edlin, it's not worth it. The editor just isn't very powerful, and you can't alter the command keystrokes.

TurboTeX. This version, from Kinch Computer, requires 3.7 Mbytes of disk space and only costs $150 for a system complete with drivers, nice utilities, and a virtual memory simulation that allows processing of larger files. I wanted to like TurboTeX, but it's very slow. I ran the first couple of chapters of my book through each TeX. TurboTeX took 10.5 minutes; the others each took less than three minutes.

TurboTeX comes with drivers for LaserJet and PostScript printers and for Epson and IBM dot-matrix printers, including the Epson LQ. You have to buy the CM files separately. It costs an additional $72 to buy all of the CM files you'll need. Again, the LaserJet driver didn't work with my compatibles. TurboTeX also comes with a complete METAFONT implementation that is separately installable. It requires 1.99 Mbytes of disk space.

For $300, you can get the source option, which also includes the executables. This TeX and METAFONT, written in C instead of Pascal, were translated using the Pascal-to-C translator that comes with the source option. TurboTeX also comes with a previewer that you can compile.

The source will compile under Microsoft's version 5.0 C compiler and under Unix or Xenix. For the Unix/Xenix version you get changes for TeX that let it process large files. For those of you who already know a lot about TeX, I'm talking about the recent 64-bit TeX changes. I haven't compiled these files, but I'll be working on a review

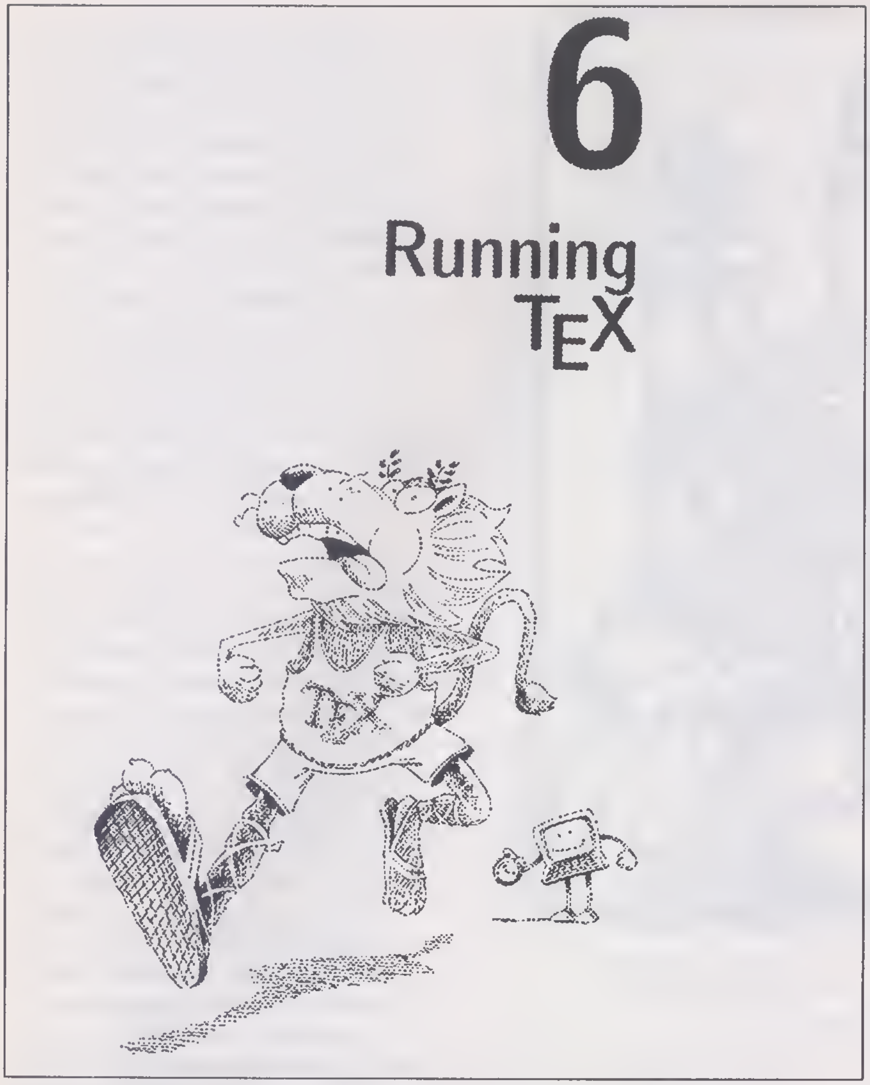

Scanner input from *The TeXbook.* (Illustration by Duane Bibby. Reprinted with permission of Addison-Wesley, Reading, Mass. ©1984 The American Mathematical Society.)

of Xenix and Unix on the 386. I'm going to use the source files to try out the compilation systems.

If TurboTeX ever lives up to its name, this could well be the version of TeX to have. Maybe in another release or two, it will. I'll let you know if it does.

***TeXT1* macro package.** This macro package, which I really like, is what I used throughout the review. It was written at Washington State University. This outstanding product sells for $150 and requires 1.9 Mbytes of disk space. *TeXT1* comes with the METAFONT input files for the font files you need, as well as the TFM files. You can save about 291 Kbytes of disk space by removing the METAFONT files.

Since you've seen throughout the review what *TeXT1* offers you, I won't spend too much time with it now. I do want to mention, however, that it's very easy to change defaults. For example, I used \subheadd{*TeXT1*} to create the current subsection title. By default, this command prints in 12-point roman italics. I was able to change it using the command

```
\subheaddformat{
 \titleformat{%
   \ninepoint \bd
     \subheaddtitle.
  }}
```

You can also control how and if sections are numbered and whether or not an entry is made in the table of contents (if one is being printed). Each of the sectioning commands has its own format command. So do the other commands that you might want to modify.

As I said before, I think this package is definitely worth having, especially if you are new to TeX.

HiJaak. This product isn't associated with TeX, and unlike the products I've reviewed up to now, you can't get this through TUG. But it's almost a must if you are going to import graphics into your TeX documents.

Available from Inset Systems, HiJaak lets you convert from one graphics format to another. For example, it lets you convert from PostScript to PCL, to TIF for scanners, to WordPerfect, and to Lotus format. While it can't change any of these to PostScript, you can convert between the rest. It requires 320 Kbytes of disk space and sells for $149.

One particularly useful feature is a switch that will keep HiJaak from putting a formfeed at the end of the file it creates. This was a last-minute add-on because of user complaints. But because it's an add-on, there are two problems with it. First, you can't use it interactively, only in batch. Second, you can't convert to the same file type, say PCL to PCL, using the switch to remove a formfeed.

I used HiJaak to produce the figures at the top of this page and the next one from TIF scanner files.

Hardware for TeX

In this part of the review, I want to tell you about the peripherals that I used. I wanted a laser printer that would produce the best TeX output possible without using a typesetter. I wanted a graphics card that would let me use the previewers most effectively, and I wanted a monitor that would really let me preview what I would get on the laser printer. I also wanted to include scanned images so that I could produce professional-looking documents. Since this review is about TeX, I'll look at the peripherals mainly from a TeX perspective; each

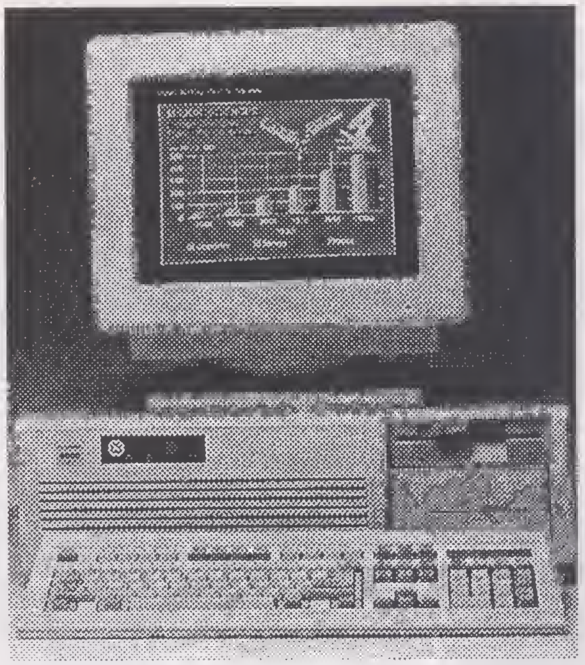

Scanner input from a computer catalog. The image was scanned at its original size, using 300 dpi.

review is just long enough to show how the peripheral enhances TeX.

ScanJet Plus. This Hewlett-Packard scanner lists for $2,190, but shopping around, you can get it for $1,560. The easy-to-use software that comes with ScanJet Plus runs under Microsoft Windows. A run-time version of Windows is included.

The figures at the top of this page and the previous page were scanned using ScanJet Plus. The Running TeX figure is an example of line art—art consisting of lines and shapes that can be drawn by hand. It was reduced to 75 percent of its original size.

Both examples were produced with the Coarse Fatting dithering pattern. Dithering is an image-processing technique used to produce continuous shades of gray. Bayer, Vertical Line, and Fatting patterns are also supported, but Coarse Fatting usually looks the best.

The figures show the quality produced by printing a scanned image at 300 dpi. If you have a 1270 or 2540 dpi Linotronic printer, or something comparable, the results will be even better. In fact, the ScanJet Plus offers 4-bit and 8-bit greyscaling. The manual that comes with the software gives an example of an 8-bit grayscale image printed at 1270 dpi. The example, an old boarded-up door, looks like a photograph. It's so clear, you can see the heads of nails that stick out.

The only problem I found with the software is that its text-recognition capability for producing ASCII files from typeset material isn't very good. In particular, I tried to scan some TeX output that contains text and Ada code. The results were totally unusable. I tried a third-party vendor for text-recognition software, but it wasn't much better.

Acer LP-76 printer. I first printed the review with this laser printer. I found the image to be crisp and clear, but the character strokes weren't dark enough. That's why I used the Brother printer instead. Aside from that, I like this printer.

The resident fonts are Courier and Lineprinter. Courier comes in 12-point, while Line Printer comes in 8.5-point. There are 23 symbol sets. Other printers have a better collection of resident fonts, but these fonts aren't important for TeX, since you get much more from the CM files. There's a slot for a font cartridge—two with an adapter. HP-GL emulation is available with a separate cartridge. The fonts, sizes, and character sets are selectable from an easy-to-use menu panel.

The default memory size is 512 Kbytes, but you can get 1-, 2-, or 4-Mbyte expansion boards. Although 512 Kbytes should be enough for your TeX needs, if you use software like Harvard Presentation Graphics or Lotus Freelance Plus to provide high-resolution output, you should go with 1 Mbyte.

An especially nice feature of this printer is the separation of the toner, cleaner, and drum. Usually they come as one replaceable unit, which costs a fair amount to replace. But in the LP-76, they can be replaced separately. The toner comes in a separate cartridge for about $20-$30, depending on where you buy, so it's a lot cheaper to operate. The printer lists for $1,995, but by shopping around you can get it for $1,565.

This printer looks almost identical to the Okidata Laserline 6, both outside and inside. The Laserline uses the same toner cartridges. As an alternative to the LP-76, you might consider the Laserline, but it only comes standard with 128 Kbytes and is only expandable to 512 Kbytes. The character strokes are a little bit thicker than they should be. Emulation of several dot matrix printers is provided by a software package. By shopping around you can get the Laserline 6 for $1,665.

Brother HL-8e. The quality of this laser printer should be apparent by looking at the review. This printer comes packed with Brougham, Anelia, and Letter Gothic fonts, as well as two Brother fonts, TMS and HLV. Brougham comes in 10, 12, and 15 points. Anelia comes in 12.5 and Letter Gothic in 8.5 points. The other two come in 10 points. In addition to LaserJet emulation, you also get HP-GL, Diablo-630, Epson FX-80, IBM ProPrinter XL, and Twinwriter WP and DP emulation. Again, you don't need all of the fonts for TeX, but they're there if you want them. Features are selected by an easy-to-use panel menu.

The Brother HL-8e comes standard with 512 Kbytes of memory, but you can extend to 1 or 2 Mbytes. There are slots for two font cartridges. The toner, cleaner, and drum come in one unit, which costs about $92 to replace. This printer lists for $2,895, but by shopping around you can get it for $1,860.

Mitsubishi monitor. The Diamond Scan 16 is an outstanding 16-inch monitor. It can be used as an analog monitor or a multisync monitor. With the right cable, you can plug it into an IBM PC or a Mac II.

This monitor lists for $2,069. By shopping around you can get it for $1,200, so it's not cheap. But I like the quality of the picture. The extra size was really useful when I used the previewers, since I could keep more of a page on the screen by using a smaller magnification.

My editor can be configured to operate at 50 lines per screen, as can the graphics adapter I used. The Diamond Scan handled this without any problems. The text is not difficult to read considering how small the characters are, but I still don't think I would want to use this mode all the time. When I need to, however, I can use the editor in this mode to scan files.

One other nice feature of this monitor is that you can adjust it to handle graphics cards of varying frequencies. The simplest way to describe it is that it's similar to tuning in a cable-ready television.

The price of this monitor may be a little steep for many of you, but if you are considering a new monitor, take a look at it. Of course, if you do, you may spend more for a monitor than you planned to.

VGA Wonder. This graphics card, from ATI Technologies, is capable of high resolutions in multiple colors. The 256-Kbyte version can produce 16 colors at a resolution of 800x600, and four colors at a resolution of 1024x768. With 512 Kbytes, it can produce 256 colors and 16 colors in the same resolutions.

The card has built-in mouse support and includes a Microsoft-compatible mouse. It has support for 132 columns by 60 lines and can operate in any video mode that your software can be configured to. There are no switches to set or jumpers to select. The card detects what monitor type you're using and what mode you're in. The 256-Kbyte version lists for $499 and the 512-Kbyte version lists for $699.

This board is particularly useful when using a previewer, since the previewers don't support higher resolution. Whenever I wanted to use Preview, I switched to Hercules 720x348 mode using the software that comes with the card. When I finished, I switched back to VGA. (Of course, I used a BAT file to do this.) For this reason, I recommend the VGA Wonder.

The only problem with this card is that you can't disable its automatic mode detection. This means that if you use a program such as VM386 or Unix with VP/ix, you'll hang the system if you switch modes. If you plan on using any program that assumes video registers are under its complete control, you'd better get a different card.

Summary

I've taken you through the world of TeX on the IBM PC and given you my recommendations on a TeX software configuration. I've also told you about the hardware that I used. With the right combination of software and hardware, using TeX on your PC will let you produce typeset-quality documents of any complexity with relative ease. And, unless you use a great deal of graphics in documents that must be turned around quickly, TeX outshines WYSIWYG systems.

Acer Technologies Corp.
401 Charcot Ave.
San Jose, CA 95131
(408) 922-0333
Reader Service 21

ArborText, Inc.
535 W. William St., Suite 300
Ann Arbor, MI 48103
(313) 996-3566
Reader Service 22

ATI Technologies, Inc.
3761 Victoria Park Ave.
Scarborough, Ontario
M1W 3S2 Canada
(416) 756-0720
Reader Service 23

Brother International Corp.
8 Corporate Place, PO Box 59
Piscataway, NJ 08855-0159
(201) 981-0300
Reader Service 24

Hewlett Packard
16399 W. Bernardo Drive
San Diego, CA 92127-1899
(619) 487-4100
Reader Service 25

Inset Systems Inc.
71 Commerce Drive
Brookfield, CT 06810
(203) 755-5866
Reader Service 26

Kinch Computer Company
501 S. Meadow St.
Ithaca, NY 14850
(607) 273-0222
Reader Service 27

Micro Publishing Systems, Inc.
1120 Hamilton Street, Suite 300
Vancouver, B. C.
V6B 2S2 Canada
(604) 687-0354
Reader Service 28

Mitsubishi Electronics America, Inc.
991 Knox Street
Torrance, CA 90502
(213) 324-6578
Reader Service 29

Personal TeX, Inc.
12 Madrona Ave.
Mill Valley, CA 94941
(415) 388-8853
Reader Service 30

WSUCSC Publications
Computing Service Center
Washington State University
Pullman, WA 99164-1220
(509) 335-0411
Reader Service 31

Zar Limited
PO Box 372
Pasadena, CA 91102
(818) 794-1224
Reader Service 32

Call for Papers

The 10th
International
Conference on

Distributed Computing Systems (ICDCS-10)

May 28-June 1, 1990 • Paris, France

CONFERENCE COMMITTEE

General Co-Chairpersons
G. Le Lann, INRIA (France)
W. Kohler, DEC (USA)

Program Committee Co-Chairpersons
R. Popescu-Zeletin, GMD-FOKUS (West Germany)
J. Stankovic, Univ. of Massachusetts/Amherst (USA)

Program Committee Vice Chairpersons

Architectures
L. Pouzin, CNET (France)

Communication Protocols
W. Bux, IBM Zurich (Switzerland)

Distributed Operating Systems
S. Mullender, CWI Amsterdam (The Netherlands)

Scheduling
T. Casavant, Purdue Univ. (USA)

Distributed Databases
W. Burkhard, Univ. of California/San Diego (USA)

Languages
E. Moss, Univ. of Massachusetts/Amherst (USA)

Tools and Methodologies
H. Weber, Univ. of Dortmund (West Germany)

Modelling and Performance Evaluation
D. Eager, Univ. of Saskatchewan (Canada)

Reliability and Fault Tolerance
H. Kopetz, Univ. of Vienna (Austria)

Distributed Algorithms and Theory
M. Herlihy, CMU and DEC Research (USA)

Real Time Systems
L. Sha, CMU and SEI (USA)

Applications
M. Sloman, Imperial College (UK)

The conference covers all aspects of Distributed Computing. The conference organizers are interested in papers dealing with innovative research work, comprehensive surveys, real implementations, or direct experience in the following areas of distributed computing:

- Architectures
- Communications Protocols
- Distributed Operating Systems
- Scheduling
- Distributed Databases
- Languages
- Tools and Methodologies
- Modelling and Performance Evaluation
- Reliability and Fault Tolerance
- Distributed Algorithms and Theory
- Real Time Systems
- Applications

Information and Instructions for Authors:

All papers are to be written in English, the language of the conference. All papers will be reviewed and those accepted will have substantive technical content. Authors of accepted papers are required to present their papers at the conference. In addition to papers, proposals for one day tutorials in any of the conference areas are solicited.

Papers must be complete and must not exceed 6000 words, where a full page figure is equivalent to 300 words. Final papers will be restricted to no more than eight IEEE model pages. Enclose a cover letter containing the following information: author's name(s), paper title, suggested primary and secondary areas from the above list, and the name, address, telephone number, and e-mail address (if available) of the author who is responsible for all correspondence and preparation of the final camera-ready version.

Please send *five* (5) copies of your complete papers with an abstract or tutorial proposals to one of the Program Co-Chairpersons:

Prof. Dr. R. Popescu-Zeletin
GMD-FOKUS
Hardenbergplatz 2
D-1000 BERLIN 12
West Germany
Phone: 49-30-25499-206

Prof. Jack Stankovic
Computer and Information Science Dept.
University of Massachusetts
Amherst, MA 01003
USA
Phone: (413) 545-0720

Submitted Paper Arrival Deadline: October 20, 1989
Acceptance Notification Date: January 22, 1990
Final Version Arrival Deadline: March 2, 1990

IEEE Computer Society
The Institute of Electrical and Electronics Engineers, Inc.

The Sixth IEEE Conference on

ARTIFICIAL INTELLIGENCE APPLICATIONS

Fess Parker's Red Lion Resort • Santa Barbara, California
March 5-9, 1990

Preliminary Call for Participation

The conference is devoted to the application of artificial intelligence techniques to real-world problems. Two kinds of papers are appropriate: Case studies of knowledge-based applications that solve significant problems and stimulate the development of useful techniques. Papers on AI techniques and principles that underlie knowledge-based systems, and in turn, enable ever more ambitious real-world applications. This conference provides a forum for such synergy between applications and AI techniques.

Papers describing significant unpublished results are solicited along three tracks:

- **Engineering/Manufacturing Track.** Contributions stemming from the general area of industrial and scientific applications.
- **Business/Decision Support Track.** Contributions stemming from the general area of business, law and various decision support applications.
- Papers in these two application tracks must: (1) Justify the use of the AI technique, based on the problem definition and an analysis of the application's requirements; (2) Explain how AI technology was used to solve a significant problem; (3) Describe the status of the implementation; (4) Evaluate both the effectiveness of the implementation and the technique used.
- **Enabling Technology Track.** Contributions focusing on techniques and principles that facilitate the development of practical knowledge based systems, and can be scaled to handle increasing problem complexity. Topics include, but not limited to: knowledge acquisition, representation, reasoning, searching, learning, software life cycle issues, consistency maintenance, verification/validation, project management, the user interface, integration, problem-solving architectures, and general tools.

Papers should be limited to 5000 words. The first page of the paper should contain the following information (where applicable) in the order shown:

- Title.
- Authors' names and affiliation. (specify student)
- Abstract: A 200 word abstract that includes a clear statement on what the original contribution is and what new lesson is imparted by the paper.
- AI topic: Knowledge acquisition, explanation, diagnosis, etc.
- Domain area: Mechanical design, factory scheduling, education, medicine, etc. Do NOT specify the track.
- Language Tool: Underlying language and knowledge engineering tools.
- Status: development and deployment status as appropriate.
- Effort: Person–years of effort put into developing the particular aspect of the project being described.
- Impact: A 20-word description of estimated or measured (specify) benefit of the application developed.

Each paper accepted for publication will be allotted seven pages in the conference proceedings. Best papers accepted in the Enabling Technology track will be considered for a special issue of IEEE Transactions on Knowledge and Data Engineering (TDKE) to appear in late 1990. Best papers accepted in the application tracks will be considered for a special issue of IEEE EXPERT, also to appear in late 1990. In addition, there will be a best student paper award of $1,500, sponsored by IBM for this conference.

In addition to papers, we will be accepting the following types of submissions:

- Proposals for Panel discussions. Topic and desired participants. Indicate the membership of the panel and whether you are interested in organizing/moderating the discussion. A panel proposal should include a 1000-word summary of the proposed subject.
- Proposals for Demonstrations. Videotape and/or description of a live presentation (not to exceed 1000 words). The demonstration should be of a particular system or technique that shows the reduction to practice of one of the conference topics. The demonstration or video tape should be not longer than 15 minutes.
- Proposals for Tutorial Presentations. Proposals of both an introductory and advanced nature are requested. Topics should relate to the management and technical development of usable and useful artificial intelligence applications. Particularly of interest are tutorials analyzing classes of applications in depth and techniques appropriate for a particular class of applications. However, all topics will be considered. Tutorials are three hours in duration; copies of slides are to be provided in advance to IEEE for reproduction.

Each tutorial proposed should include the following:

- Detailed topic list and extended abstract (about 3 pages)
- Tutorial level: introductory, intermediate, or advanced
- Prerequisite reading for intermediate and advanced tutorials
- Short professional vita including presenter's experience in lectures and tutorials.

- Proposals for Vendor Presentations. A separate session will be held where vendors will have the opportunity to give an overview of their AI-based software products and services.

IMPORTANT DATES

- September 29, 1989: Six copies of Papers, and four copies of all the proposals are due. Submissions not received by that date will be returned unopened. Electronically transmitted materials will not be accepted.
- October 30, 1989: Author notifications mailed.
- December 12, 1989: Accepted papers due to IEEE Computer Society. Accepted tutorial notes due to Tutorial Chair, Donald Kosy
- March 5-6, 1990: Tutorials
- March 7-9, 1990: Conference

Submit Papers and Other Materials to:

Se June Hong (Room 31-206)
IBM T.J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598
USA

Phone: (914) 945-2265
CSNET: HONG@IBM.COM
FAX: (914) 945-2141
TELEX: 910-240-0632

Submit Tutorial Proposals to:

Donald Kosy
Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

Phone: (412) 268-8814
ARPANET:kosy@cs.cmu.edu

Additional Information

For registration and additional conference information, contact:

CAIA-90
IEEE Computer Society
1730 Massachusetts Avenue, NW
Washington, DC 20036-1903

Phone: 202-371-1013

IEEE COMPUTER SOCIETY

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

NEW PRODUCTS

Contact or send releases to Nancy Hays, *Computer*, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720; Compmail+, n.hays

Simulation accelerator supports millions-of-gates designs

Zycad's System Evaluator logic-simulation accelerator supports system designs with up to four million modeling elements. The System Evaluator reportedly performs up to 320 million events per second (MEPS) using zero and unit delay models.

The SE features a parallel architecture in which each of up to 16 simulation modules handles 256,000 gates and contains four gate pipelines and two memory model engines. Each pipeline handles 64,000 models and includes functions needed to execute standard models in zero or unit delay. Each memory model engine supports a capacity of 16 Mbytes.

A time queue processor handles time delays of greater than one, while a wired function processor calculates all node resolution functions resulting from wired or bus models.

The SE features clamp/unclamp capability, partial network load, and truth tables programmable by individual pipelines.

The SE has standard connections to Sun, Apollo, DEC VAX, and IBM host computers, and runs the company's Zilos operating system. It comes in three models.

The SE1064 has a two-million-gate capacity, 32 Mbytes of memory, and 160 MEPS performance for $3,320,000. The SE1096 has a three-million-gate capacity, 48 Mbytes of memory, and 240 MEPS performance for $4,620,000. The SE1128 supports four million gates, 64 Mbytes of memory, and 320 MEPS speed for $5,920,000.

Shipments are scheduled for the first quarter of 1990.

Reader Service 40

Entry-level HP-PA out

Hewlett-Packard has added an entry-level minicomputer to its family of HP Precision Architecture systems. The new multiuser computer is based on the Unix operating system and the company's reduced instruction set computing architecture.

The HP 9000 Model 815S includes two serial ports, 8 Mbytes of memory, and a 16-user license for the HP-UX operating system, which follows AT&T's Unix System V Interface Definition Issue 2. A standard system costs $14,900.

A complete, bundled system for 16 users costs $29,500. This includes 8 Mbytes of memory, a 300-Mbyte disk drive, a tape cartridge system, 18 serial ports, and a 16-user license for HP-UX, which comes preloaded.

Reader Service 42

Accelerator tops Mach family

Zycad has topped its Mach family of logic and fault simulation accelerators with the Mach 2000, which supports designs with 64,000 to 1 million modeling elements.

The Mach 2000 combines a special-purpose hardware accelerator with a Sun-4 workstation, which uses a Sparc RISC processor and the Unix operating system. Local memory ranges from 32 to 128 Mbytes and disk memory from 760 to 3,040 Mbytes.

According to the company, it has developed a new algorithm that allows the Mach 2000 to carry out parallel concurrent fault simulation. All Mach models will come with the algorithm beginning in August.

Scheduled for shipment in October 1989, the Mach 2000 will cost $185,000 with logic simulation only and $240,000 with logic and fault simulation. Mach 1000 and 1500 models can be upgraded in the field for $130,000.

Reader Service 41

Zycad's Mach 2000 accelerator supports concurrent logic and fault simulation for designs with up to one million gates.

Apple opens Macintosh communications

Apple Computer now offers a range of communications products for the Macintosh PC. The new products allow users to communicate with other Macs and with other computer environments. They include enhancements to the AppleTalk network system, integration products for other computing environments, and networking and communications tools.

According to the company, AppleTalk Phase 2 is a revised version of the network system that provides addressing up to 16 million nodes and improved routing and broadcasting performance. Products include AppleTalk Internet Router ($399), TokenTalk Version 2.0 (included in TokenTalk NB Card), EtherTalk Version 2.0 (included in EtherTalk NB Card), AppleShare PC Version 2.0 ($149), LocalTalk PC Card bundle ($299), and EtherTalk for A/UX (upgrade, available in the fourth quarter). All but the last will be available in the third quarter.

Integration products include TokenTalk NB Card with TokenTalk Version 2.0 software and SMB File Transfer software ($1,250 in the third quarter), Apple Coax/Twinax Card with MacDFT software ($1,495 in the third quarter), Apple Serial NB Card ($1,195), MacDFT ($245), MacAPPC Version 1.1 (license), MacX ($299), and MacX25 ($800). All but the first two products will be available in the fourth quarter.

Networking and communications products include CL/1 servers for MVS/TSO (DB2) and VM/CMS (SQL/DS), available in the first quarter of 1990 for license fees; Macintosh Communications Toolbox (license in the third quarter); and Apple Data Modem 2400 ($499 in the third quarter).

Contact the company for detailed information on the new communications products.

AppleTalk: **Reader Service 43**
Integration: **Reader Service 43**
Networking: **Reader Service 45**

Transputers combine in SuperCluster

Paracom has released the SuperCluster parallel processing supercomputer, based on the Inmos Transputer chip. The system employs a hierarchic cluster architecture.

The Transputers include 4 Kbytes of internal static RAM, an on-chip floating-point processor, and four 20-Mbit/s links. Each processor node has up to 4 Mbytes of dynamic RAM with error detection and correction.

According to the company, the basic subunit of the supercomputer consists of the computing cluster, designed around 16 T800 Transputers operating at 25 MHz, and a network configuration unit. The NCU controls up to 96 serial communication links and can establish any processor topology within the computing cluster while reserving 32 links for connection to other cluster units.

The 64-processor SuperCluster system consists of four computing clusters, a system services cluster, and two NCU's. Additional units can be connected to form systems of 128, 256, 512, or more processors.

Software support includes the Helios operating system, MultiTool software development environment, and compilers for Occam, C, Fortran, Pascal, Prolog, ParC, and others.

Prices for a 64-node SuperCluster start at $336,500.

Reader Service 46

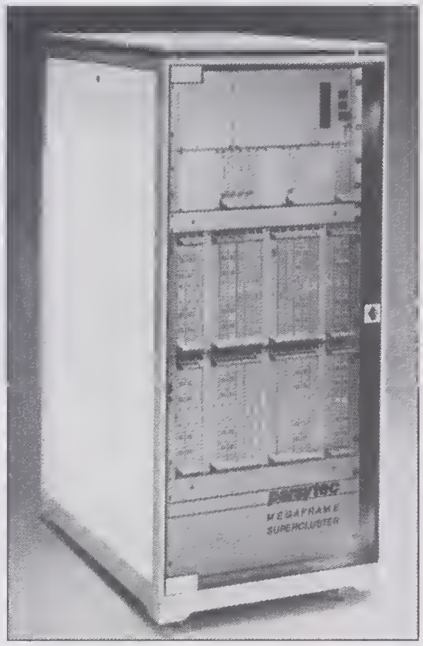

Paracom's SuperCluster combines Inmos Transputer chips in a hierarchic cluster architecture.

CAD for Mac

Schlumberger CAD/CAM has added the MacBravo Electronics Designer to its family of CAD software for the Apple Macintosh computer. According to the company, the new software combines a graphics interface with existing simulation tools.

Waveform editing tools include operations such as analog to digital, digital to analog, cut and paste, and functional computation on waves. Engineers can display, edit, or use simulation results again graphically. Electronics Designer also provides a schematics capture function.

Key features include unlimited hierarchy, unlimited number of sheets, and a graphics preview window. Options include libraries of components, such as Texas Instruments, Fairchild, RCA, Motorola, and Intel. A variety of tools allow users to create drawings that conform to ANSI, ISO, DIN, or other standards.

Like other MacBravo products, Electronics Designer contains a customization capability called the Flexible Interface Tool, Mousestrokes, and a Hypercard help facility.

Electronics Designer is scheduled for availability in the fourth quarter. It costs $1,950. The component library costs $495.

Reader Service 47

PDQ cuts timing problems

Quad Design Technology offers a component-placement program to calculate signal interconnect delays in circuit design. PDQ (Preroute Delay Quantifier) reportedly helps designers identify and eliminate system-timing problems early in the design process. The software automatically determines the optimal signal paths between components to minimize those delays.

According to the company, PDQ employs a distributed-parameter routing model to produce estimates of interconnect timing. By addressing timing at the component-placement stage, PDQ helps identify problems before physical prototyping and manufacturing.

PDQ works alone or in conjunction with Motive (Modular Timing Verifier) and TLC (Transmission Line Calculator) software. It currently runs on Sun and Apollo workstations. Prices start at $3,900.

Reader Service 48

Transputer for Apollo

Cresco Data offers the CD-TB40 transputer board and CD-TSE transputer software environment for Apollo DN3000 and DN4000 workstations and compatibles.

A single CD-TB40 board includes four 32-bit Inmos Transputers with 10 MIPS performance and a total of 7-28 Mbytes of fault-tolerant RAM. The board connects to the workstation via a direct-memory-access interface operating at up to 600 Kbytes/s.

The CD-TSE software features a remote-procedure-call interface. It reportedly allows users to integrate transputer code into new or existing Apollo programs written in Fortran, C, or Pascal.

Three CD-TB40 boards, each with its own interface, can work independently to provide up to 120 MIPS, or boards can be linked to form arrays requiring a single interface to the host.

The CD-TB40 board is compatible with the Inmos B008 Transputer evaluation board. The CD-TSE software supports all functions of the Inmos TDS2 Transputer development system.

Contact the Danish manufacturers for more information on pricing and availability.

Reader Service 49

Logic synthesis for VHDL

Synopsys claims that its HDL Compiler, used with its Design Compiler logic synthesis and optimization software, automatically translates a register-transfer-level VHDL (IEEE 1076) circuit description into gates. The Design Compiler then optimizes the circuit for speed or area based on engineer-specified constraints.

The HDL Compiler reportedly extracts a technology-independent register-transfer-level VHDL state-machine and VHDL dataflow description from an existing circuit's netlist for documentation and resynthesis. This means that the Design Compiler can synthesize redesigns into the ASIC libraries of different vendors.

The HDL Compiler with VHDL support, including extraction software, will be available in December 1989 for $22,500 and up.

The HDL Compiler was first introduced in September of 1988. It runs on Apollo, Sun, and DEC workstations.

Reader Service 50

CAE software predicts crosstalk effects

Quad Design Technology claims that its XTK (Crosstalk Tool Kit) predicts the effects of signal coupling, or crosstalk, on system performance early in the development of ICs and printed-circuit boards. According to the company, XTK models linear and nonlinear effects associated with drivers and receivers for TTL, CMOS, and ECL logic families.

XTK reportedly calculates the degree of coupling that can be expected to occur between signals in high-density circuits, and it simulates the effects of this interference on signal waveforms in a circuit or a complete system.

XTK consists of two programs: XFX (Crosstalk Field Extractor) and XNS (Crosstalk Network Simulator). The programs work separately or together. XFX uses conductor-geometry data to calculate signal-line characteristics for traces in a given configuration. It also determines the degree of coupling that will affect each set of parallel traces in a circuit network. XNS uses this information to simulate the network waveforms. Users can enter new line characteristics into XNS until "what if" analyses show reduced crosstalk and acceptable waveforms.

XTK runs on Sun workstations. Prices start at $25,000, with availability scheduled for the third quarter. Sold separately, XFX costs $15,000 and XNS costs $17,000.

Reader Service 51

DEC upgrades graphics software support

Digital Equipment now offers GKS (Graphical Kernel System) and PHIGS (Programmer's Hierarchical Interactive Graphics System) on Ultrix and VMS systems running the DECwindows environment.

DEC GKS V4 and DEC PHIGS V2 support VAX workstations, including the VAXstation 3520 and 3540, which offer 2D and 3D graphics. The 2D support is based on the X Window system, while the 3D support is based on the proposed PEX (PHIGS+ extensions for X11) standard.

A DEC PHIGS development license costs $3,500; the runtime license costs $700 for workstations. The development license for DEC GKS costs $900, while the runtime license costs $180 for all workstations.

GKS: **Reader Service 52**
PHIGS: **Reader Service 53**

Software promotes data visualization

Precision Visuals offers a software package for exploring, analyzing, and visualizing scientific data interactively. PV-WAVE (Precision Visuals' Workstation Analysis and Visualization Environment) reportedly allows scientists to select and analyze key features and trends in large data sets, then translate results into charts, graphs, contour maps, surface plots, and images.

According to the company, users can combine traditional computational analysis and graphics with image-processing techniques to look at their data in nontraditional ways, such as by combining 3D shaded surfaces with contour maps and pixel images in one graphics window. The software can display four-dimensional data in the same picture or up to six dimensions in the same window. It also creates animations of data sets by rapidly redisplaying predefined pixel images.

PV-WAVE works on Sun and DEC workstations and in VAX/VMS environments. It features direct end-user access to windows, mouse input, menus, and operating system function. It can output graphics to both gray-scale and color PostScript printers.

The company asserts that users can directly access existing applications software or subroutine libraries from within the PV-WAVE data visualization environment. This avoids the need to write new code, compile, or link.

Prices for PV-WAVE start at $3,350 for media and documentation.

Reader Service 54

Stand-alone workstations for scientific visualization

Tom Culviner, Staff Editor, IEEE Computer Graphics and Applications

Scientific visualization is performed on many different systems and configurations, ranging from Macintosh IIs (using data from larger computers or supercomputers) to specialized visualization superworkstations and graphics supercomputers. This roundup consists of products that are used as or could be used as stand-alone visualization workstations. Some of the high-end machines listed here seem to transcend the category of workstations, both in their prices and in their multiuser options. They have been included because the multiuser option can be adapted for multiple-monitor operation by a single user.

All the machines here have specialized graphics hardware — a graphics subsystem, a graphics accelerator, or one or more graphics processors. Most have display resolution of 1,280 × 1,024 pixels and 19-inch color displays with at least 256 colors displayable from a 16.7-million-color palette. Many offer 16.7 million displayable colors.

When Unix is given as the operating system, often what the machine runs is the manufacturer's version of Unix. For example, Alliant's operating system is Concentrix, and Silicon Graphics' is Irix; both are Unix-based systems.

Anyone who has read literature from vendors has seen the numerous figures — MIPS, Mflops, and vectors and pixels per second — with which vendors describe their products' performance. Such figures are missing from the table below. Vendors use their own programs and models to determine the CPU and graphics performance of their machines, so comparing figures given by one company with figures from another company will not result in a valid comparison of performance.

Company	CPU	CPU Memory in Mbytes	Floating-point Performance in Mflops (see text above)	Operating System	Monitor Resolution in Pixels
Alliant VFX/40	1-4 64-bit processors	32-128	2.4 (1 processor)	Unix	1,280 × 1,024
Alliant VFX/80	1-8 64-bit processors	32-224	2.4 (1 processor)	Unix	1,280 × 1,024
Apollo 10000VS	1-4 64-bit processors	8-128	5.1 (1 processor)	Unix	1,280 × 1,024
Ardent Titan	1-4 64-bit processors	16-128	6.5 (1 processor)	Unix	1,280 × 1,024
DEC VAXstation 3520	2 32-bit processors	8-64	NA	VMS or Unix	1,280 × 1,024
DEC VAXstation 3540	4 32-bit processors	8-48	NA	VMS	1,280 × 1,024
Hewlett-Packard 9000					
825SRX	1 32-bit processor	8-96	0.53	Unix	1,280 × 1,024
835SRX	1 32-bit processor	8-96	1.8	Unix	1,280 × 1,024
835 TurboSRX	1 32-bit processor	8-96	1.8	Unix	1,280 × 1,024
Silicon Graphics Power Iris					
4D/120GTX	2 32-bit processors	8-128	NA	Unix	1,280 × 1,024
4D/220GTX	2 32-bit processors	8-128	NA	Unix	1,280 × 1,024
4D/240GTX	4 32-bit processors	8-128	NA	Unix	1,280 × 1,024
Stellar GS1000	4 64-bit processors	16-128	9.8	Unix	1,280 × 1,024
Stellar GS2000	4 64-bit processors	16-128	NA	Unix	1,280 × 1,024
Stellar GS2500	5 64-bit processors	16-128	NA	Unix	1,280 × 1,024
Sun Sparcstation 370 with TAAC-1	1 32-bit processor	8-56	NA	Unix	1,152 × 900 1,024 × 1,024

Jack J. Dongarra of Argonne National Laboratory in Argonne, Illinois, publishes a report entitled ''The Performance of Various Computers Using Standard Linear Equations Software in a Fortran Environment.'' The single performance figure given for machines in the table below comes from the June 6, 1989, issue of that report (''NA'' indicates that the machine is not listed in the report). In his report, Dongarra emphasizes that his results ''reflect only one problem area: solving dense systems of equations using the Linpack programs in a Fortran environment.'' However, the figures are a good starting point for comparison, like the EPA figures on gas mileage for cars. The figures cited in the table below come from the listing for full-precision Linpack performance using all Fortran. (Some vendors provide in their product literature figures from Dongarra's full-precision all-Fortran list; other vendors provide figures from other tables in the report where the figures are higher, or use their own figures.)

Unfortunately, there is no such standard measure of graphics performance, a lack that the graphics community has recognized and is trying to remedy. Because the vendor-supplied figures for graphics performance cannot be used for controlled comparison, they are not given here.

The column labeled "Comments" does not give exhaustive comments on the machines; it simply reflects some of what vendor representatives and literature emphasized as important features.

The base price supplied by the company is usually for a one-processor configuration with the minimum CPU memory, with color capabilities and some disk capacity and sometimes tape capacity. Wide ranges of options in disk and tape storage are available for added cost.

''People buy cars, not engines,'' David Salzman and Jack Grimes note in the July 1989 issue of *IEEE Computer Graphics and Applications* (p. 28), as an analogy for how people go about buying graphics superworkstations. Because each user has unique needs and problems to solve, the most the performance figures and sketches below can do is give some preliminary information on where to look for a solution.

Colors Displayed/ Colors Available	Comments	Base Price	Contact Phone Number	Reader Service Number
16.7M/16.7M	Configurable with up to 8 independent graphics processors; 1-4 users	$200,000	(800) 622-1113; (800) 922-2405 in Mass.	60
16.7M/16.7M	Configurable with up to 16 independent graphics processors; 1-8 users	$359,000	(800) 622-1113; (800) 922-2405 in Mass.	60
16.7M/16.7M	Two-board graphics subsystem; 3D RISC drawing engine; does not require use of proprietary graphics interface	$94,900	(508) 256-6600	61
16.7M/16.7M	Accepts Cray and VAX source code; bundled with visualization environment and Dore graphics library	$79,000	(408) 732-0400	62
256/256	Graphics accelerator; accepts PHIGS; 16.7 M color upgrade available	$35,900	(508) 897-5111	63
256/256		$55,900	(508) 897-5111	63
256/16.7M	Graphics subsystem; RISC architecture; lower end field upgradable to higher end; 16.7M displayable colors option available	$56,500	(800) 752-0900	64
256/16.7M		$69,500	(800) 752-0900	64
256/16.7M		$63,500	(800) 752-0900	64
16.7M/16.7M	Graphics subsystem; RISC architecture; board-swapping upgrade path; video digitizer board available	$94,900	(415) 960-1980	65
16.7M/16.7M		$114,900	(415) 960-1980	65
16.7M/16.7M		$139,900	(415) 960-1980	65
16.7M/16.7M	Comes with Application Visualization System software	$90,000	(617) 964-1000	66
16.7M/16.7M	Comes with Application Visualization System software; dual-user option available	$125,000	(617) 964-1000	66
16.7M/16.7M		$140,000	(617) 964-1000	66
16.7M/16.7M	Video from TAAC-1 accelerator board can be displayed in a window or on a separate monitor	$78,900	(919) 469-8300	67

ADVERTISER INDEX

Academic Press, Inc. . . . 8
Agouron Pharmaceuticals, Inc. . . . 142
Alsys Inc. . . . 9
CACI Products Company . . . 1, Cover IV
CAIA '90 . . . 122
Computer Arithmetic Symposium . . . 131
Computer Systems Architects . . . Cover II
GTE Government Systems Corp. . . . 107
The Heinz-Nixdorf Institute . . . 51
IBM Corp., Data Systems Division . . . 143
ICDCS-10 . . . 121
IEEE Computer Society Membership . . . 139, 152
IEEE Computer Society Ombudsman . . . 142
IEEE Computer Society Press . . . 146
Interactive Software Engineering Inc. . . . 6-7
Local Computer Networks Conference . . . 58-59
Meta Software . . . 45, 52
MetaWare . . . 66
NYNEX Science & Technology Center . . . 67
Oakland Group, Inc. . . . 37
Office Information Systems Conference . . . 106
Reliable Distributed Systems Symposium . . . 102-103
Software Productivity Consortium . . . 26
StarSys, Inc. . . . 38
Zortech Inc. . . . 5
Classified Advertising . . . 144-145

FOR DISPLAY ADVERTISING INFORMATION, CONTACT:

Northern California and Pacific Northwest: Roy McDonald Assoc. Inc., 5915 Hollis St., Emeryville, CA 94608; (415) 653-2122.
Jim Olsen, P.O. Box 696, Hillsboro, OR 97123; (503) 640-2011.
Southern California and Mountain States: Richard C. Faust Co., 24050 Madison St., Suite 100, Torrance, CA 90505; (213) 373-9604.
Southwest: The House Co., 5252 Westchester, Suite 280, Houston, TX 77005; (713) 668-1007.
Midwest: The Kingwill Company, 4433 W. Touhy Ave., #540, Lincolnwood, IL 60646; (312) 675-5755.
East Coast: Atlantic Representative Group, 349 Maple Pl., Keyport, NJ 07735; (201) 739-1444.
New England: Arpin Associates, P.O. Box 6444, Holliston, MA 01746; (508) 429-8907.
Europe: Heinz J. Görgens, Parkstrasse 8a, D-4054 Nettetal 1 - Hinsbeck (F.R.G.); phone: (0 21 53) 8 99 88; telex 841(17)2153310=HJG tlx d.

For production information, conference, and classified advertising, contact Heidi Rex or Marian Tibayan.

COMPUTER, 10662 Los Vaqueros Cir., Los Alamitos, CA 90720; phone (714) 821-8380; fax (714) 821-4010.

PRODUCT INDEX

	RS#	Page#
BOARDS		
Transputer board	49	125
CIRCUITS		
Graphics card	23	120
IC announcements	120-132	129
COMPONENTS		
Microsystem announcements	135-147	130
CONFERENCES & COURSES		
CAIA '90	—	122
Computer Arithmetic	—	131
ICDCS-10	—	121
Local Computer Networks	—	58-59
Office Information Systems	—	106
Reliable Distributed Systems	—	102-103
DATA ACQUISITION		
Graphics converter	26	118
Logic-simulation accelerator	40-41	123
DATA COMMUNICATIONS		
Communications product	43-45	124
Logic-simulation accelerator system	40	123
I/O RELATED EQUIPMENT		
Laser printer	21, 24	119
Monitor	29	120
Scanner	25	119
PUBLISHERS		
CS membership	—	139, 152
CS ombudsman	—	142
Scientific Computing	4	8
Visualization in Scientific Computing	—	146
TEX Made Easy	32	115
RECRUITMENT		
Classified advertising	—	144-145
Technical professional		26, 51, 67, 107, 142, 143
SOFTWARE		
Ada compiler	—	9
CAD/CAE software	47, 51	124, 125
C compiler	8	66
C tool	2	5
Data visualization	54	125
Graphics software	52-53	125
Logic synthesis software	50	125
Simulation package	—	1, C.IV
SYSTEMS		
Computer	42, 46	123, 124
Design tool	7	45
Graphics development software	9	52
Interface management system	5	37
Object-oriented environment	3	6-7
Productivity tool	6	38
Transputer	1	C.II
Typesetting system	22, 27-28, 30-31	116-118
TEST & MEASURING EQUIPMENT		
Timing verifier	48	124

READER SERVICE CARD

COMPUTER

August 1989 issue (Card void after January 1990)
Please print or use peel-off address label.

Name ____________________

Title ____________________

Company ____________________

Address ____________________

City ____________________

State/ZIP ____________________

Country ____________ Phone ____________

Please send (circle those you want)

201 Publications catalog	**203** Student membership information
202 Membership information	**204** Application for senior membership

Reader interest
(Circle what you liked;
add comments on the back)

- **A** Guest Ed's Intro
- **B** Expanding Opportunities
- **C** Fluid Flow Data Sets
- **D** Turbulent Flow/Flames
- **E** 3D Medical Data
- **F** Earth Sciences
- **G** Semiconductors
- **H** Research Overviews
- **J** President's Message
- **K** Letters
- **L** Open Channel
- **M** Standards
- **N** CS News/Update
- **P** Product Reviews
- **Q** New Products
- **R** ICs/Microsystems
- **S** Conferences
- **T** Calendar/Call
- **U** Book Reviews
- **U** New Lit/CS Magazines

Product inquiries
(circle numbers for products and advertisements
on which you want more information)

1

1	21	41	61	81	101	121	141	161	181
2	22	42	62	82	102	122	142	162	182
3	23	43	63	83	103	123	143	163	183
4	24	44	64	84	104	124	144	164	184
5	25	45	65	85	105	125	145	165	185
6	26	46	66	86	106	126	146	166	186
7	27	47	67	87	107	127	147	167	187
8	28	48	68	88	108	128	148	168	188
9	29	49	69	89	109	129	149	169	189
10	30	50	70	90	110	130	150	170	190
11	31	51	71	91	111	131	151	171	191
12	32	52	72	92	112	132	152	172	192
13	33	53	73	93	113	133	153	173	193
14	34	54	74	94	114	134	154	174	194
15	35	55	75	95	115	135	155	175	195
16	36	56	76	96	116	136	156	176	196
17	37	57	77	97	117	137	157	177	197
18	38	58	78	98	118	138	158	178	198
19	39	59	79	99	119	139	159	179	199
20	40	60	80	100	120	140	160	180	200

READER SERVICE CARD

COMPUTER

August 1989 issue (Card void after January 1990)
Please print or use peel-off address label.

Name ____________________

Title ____________________

Company ____________________

Address ____________________

City ____________________

State/ZIP ____________________

Country ____________ Phone ____________

Please send (circle those you want)

201 Publications catalog	**203** Student membership information
202 Membership information	**204** Application for senior membership

Reader interest
(Circle what you liked;
add comments on the back)

- **A** Guest Ed's Intro
- **B** Expanding Opportunities
- **C** Fluid Flow Data Sets
- **D** Turbulent Flow/Flames
- **E** 3D Medical Data
- **F** Earth Sciences
- **G** Semiconductors
- **H** Research Overviews
- **J** President's Message
- **K** Letters
- **L** Open Channel
- **M** Standards
- **N** CS News/Update
- **P** Product Reviews
- **Q** New Products
- **R** ICs/Microsystems
- **S** Conferences
- **T** Calendar/Call
- **U** Book Reviews
- **U** New Lit/CS Magazines

Product inquiries
(circle numbers for products and advertisements
on which you want more information)

2

1	21	41	61	81	101	121	141	161	181
2	22	42	62	82	102	122	142	162	182
3	23	43	63	83	103	123	143	163	183
4	24	44	64	84	104	124	144	164	184
5	25	45	65	85	105	125	145	165	185
6	26	46	66	86	106	126	146	166	186
7	27	47	67	87	107	127	147	167	187
8	28	48	68	88	108	128	148	168	188
9	29	49	69	89	109	129	149	169	189
10	30	50	70	90	110	130	150	170	190
11	31	51	71	91	111	131	151	171	191
12	32	52	72	92	112	132	152	172	192
13	33	53	73	93	113	133	153	173	193
14	34	54	74	94	114	134	154	174	194
15	35	55	75	95	115	135	155	175	195
16	36	56	76	96	116	136	156	176	196
17	37	57	77	97	117	137	157	177	197
18	38	58	78	98	118	138	158	178	198
19	39	59	79	99	119	139	159	179	199
20	40	60	80	100	120	140	160	180	200

READER SERVICE CARD

COMPUTER

August 1989 issue (Card void after January 1990)
Please print or use peel-off address label.

Name ____________________

Title ____________________

Company ____________________

Address ____________________

City ____________________

State/ZIP ____________________

Country ____________ Phone ____________

Please send (circle those you want)

201 Publications catalog	**203** Student membership information
202 Membership information	**204** Application for senior membership

Reader interest
(Circle what you liked;
add comments on the back)

- **A** Guest Ed's Intro
- **B** Expanding Opportunities
- **C** Fluid Flow Data Sets
- **D** Turbulent Flow/Flames
- **E** 3D Medical Data
- **F** Earth Sciences
- **G** Semiconductors
- **H** Research Overviews
- **J** President's Message
- **K** Letters
- **L** Open Channel
- **M** Standards
- **N** CS News/Update
- **P** Product Reviews
- **Q** New Products
- **R** ICs/Microsystems
- **S** Conferences
- **T** Calendar/Call
- **U** Book Reviews
- **U** New Lit/CS Magazines

Product inquiries
(circle numbers for products and advertisements
on which you want more information)

3

1	21	41	61	81	101	121	141	161	181
2	22	42	62	82	102	122	142	162	182
3	23	43	63	83	103	123	143	163	183
4	24	44	64	84	104	124	144	164	184
5	25	45	65	85	105	125	145	165	185
6	26	46	66	86	106	126	146	166	186
7	27	47	67	87	107	127	147	167	187
8	28	48	68	88	108	128	148	168	188
9	29	49	69	89	109	129	149	169	189
10	30	50	70	90	110	130	150	170	190
11	31	51	71	91	111	131	151	171	191
12	32	52	72	92	112	132	152	172	192
13	33	53	73	93	113	133	153	173	193
14	34	54	74	94	114	134	154	174	194
15	35	55	75	95	115	135	155	175	195
16	36	56	76	96	116	136	156	176	196
17	37	57	77	97	117	137	157	177	197
18	38	58	78	98	118	138	158	178	198
19	39	59	79	99	119	139	159	179	199
20	40	60	80	100	120	140	160	180	200

Editorial comments

I liked: ______________________

____

____

____

I disliked: ______________________

____

____

I would like to see: ______________________

____

____

____

For reader-service inquiries, see other side.

PO Box is for reader-service cards only.

PLACE
POSTAGE
HERE

COMPUTER
Reader Service Inquiries
PO Box 16508
North Hollywood, CA 91615-6508
USA

Editorial comments

I liked: ______________________

____

____

____

I disliked: ______________________

____

____

I would like to see: ______________________

____

____

____

For reader-service inquiries, see other side.

PO Box is for reader-service cards only.

PLACE
POSTAGE
HERE

COMPUTER
Reader Service Inquiries
PO Box 16508
North Hollywood, CA 91615-6508
USA

Editorial comments

I liked: ______________________

____

____

____

I disliked: ______________________

____

____

I would like to see: ______________________

____

____

____

For reader-service inquiries, see other side.

PO Box is for reader-service cards only.

PLACE
POSTAGE
HERE

COMPUTER
Reader Service Inquiries
PO Box 16508
North Hollywood, CA 91615-6508
USA

IC Announcements

Company, Model, Function	Comments	R.S. No.
Analog Devices AD1679, AD1779 Sampling ADCs	Two 14-bit monolithic sampling analog-to-digital converters with recursive subranging for a 10-µs total conversion time and a 100 kilosamples/s throughput rate. AD1679 is for 8-bit buses; AD1779 is for 16-bit and wider buses. Available in two grades in a 28-pin ceramic or plastic DIP. Cost: $39 for JN grade, $43 for KN grade.	120
Chips and Technologies 82C480 Graphics processor	A graphics processor having register-level compatibility with IBM's 8514/A graphics standard. Supports AT and PS/2 buses. Extends the IBM 8514/A adapter capability to noninterlaced 1,024×768 display resolutions. Comes in a 160-pin plastic flatpack. Samples in September. Cost (100s): $99.	121
Cirrus Logic LC-GP425 RPX	A raster printer accelerator featuring trapezoid fill and parallel processing capabilities. Provides direct-memory-access control logic and a bit manipulation engine that off-loads the system microprocessor during pattern fill operations. Cost (1,000s): $35.	122
Gould AMI 631000, 631001 ROMs	One-megabit ROMs featuring 150- or 200-ns response times. The 631000 series ROMs require a 5V power supply and come in 28-pin DIPs. The 32-pin 631001 ROMs come in surface-mount PLCCs. Both come in commercial, industrial, and military temperature-range versions. Cost (10,000s): starts at $5.50.	123
Gould AMI 22CV10, 20CG10 PLDs	Programmable electrically erasable logic devices operating at 25 ns and consuming 55 mA (22CV10) and 45 mA (20CG10) in standby mode. Up to 12 independently programmable I/O macro cells. Come in 24-pin, windowless, 300-mil DIPs. Cost (1,000s): $7 for 22CV10; $4.50 for 20CG10.	124
Inmos IMS A121 Image processor	A 2D discrete cosine transform image processor. Enables image data compression and decompression. Does an 8×8 matrix multiplication between an 8×8 set of data values and an 8×8 set of coefficients. Comes in 44-pin PLCCs. Cost (1,000s): $80.	125
Integrated Device Technology IDT7050, IDT7052 SRAMs	Four-port SRAMs organized 1K×8 (7050) and 2K×8 (7052). Each port features independent asynchronous address, data, and control lines and communicates via a standard SRAM interface. Available in 108-pin plastic and ceramic PGAs and 132-pin plastic and ceramic quad flatpacks. Cost (100s): $212.50 for 7050, $297.50 for 7052 in 25-ns ceramic PGAs.	126
Intel 27C240 EPROM	A 4-Mbit EPROM organized 256K×16 capable of storing more than 4 million bits. A pin-compatible upgrade from Intel's 1-Mbit 27C210. Comes in a JEDEC-standard, 40-pin ceramic DIP in both 150- and 200-ns access versions. Cost (10,000s): $100 for 200-ns version.	127
Signetics 87C752 Microcontroller	A 28-pin microcontroller in the 80C51 family, with analog-to-digital converter and pulse-width modulator functions. Software compatible with the 80C51. Comes in DIP and PLCC packages. Cost: $3 for ROM mask version, $10 for a one-time programmable EPROM version, $35 for a user-erasable quartz-windowed EPROM version.	128
Teledyne Semiconductor TSC820 Measurement system	A measurement-system chip containing an ADC, a logic probe, a frequency counter, and an LCD driver. Also includes decimal point drivers, a low-battery detector, a buzzer-output driver, and peak reading hold. Has a resolution of 3.75 digits. Comes in 40-pin DIP and 44-pin PLCC and plastic flatpack packages. Cost (100s): starts at $6.	129
Teledyne Semiconductor TSC835 ADC	A precision analog-to-digital converter for data acquisition boards that interface with IBM PC and Apple Macintosh computers. Features 0.005% resolution, 5 conversions/s, and 100 µV input sensitivity. Interfaces with a UART for serial data transmission. Comes in a variety of DIP and surface-mount packages. Cost (1,000s): starts at $4.48	130
Texas Instruments TLC32044 Analog interface	An analog-interface chip with flexible filter bandwidths, one of a series of universal DSP peripherals. Features a nominal bandwidth of 100-3,800 Hertz. Measured delays of 300 and 250 µs. Dissipates 385 mW max. Comes in a 28-pin plastic DIP. Two grades. Cost (1,000s): $19.45 for CN grade, $22.22 for IN grade.	131
Western Digital Imaging PWGA1 Graphics chip set	A graphics chip set called Personal Workstation Graphics Array 1. Compatible with IBM's 8514/A Display Adapter. Has a 4-bit bitblt mode and 60-MHz clock speed. Supports interlaced and noninterlaced monitors up to 1,280×1,024 resolution. Production in third quarter. *Cost (100s): $142.*	132

Microsystem Announcements

Company, Model, Function	Comments	R.S. No.
Arcom Control Systems SG84X Genlock board	A single-Eurocard genlock/graphics board. Works with Arcom's ARCTC-based SG84 STEbus graphics controller board. Operates in two modes, to superimpose text on video or mix graphics and video signals. Compatible with PAL gray-scale composite video. Cost: £142; £485 for SG84.	135
Asante Technologies MacCon II/TR Interface card	A Token-Ring network interface card for Apple Macintosh computers. Complies with the IEEE 802.5 Token-Passing Protocol Standard. Features a multicolored diagnostics LED. Optional Logical Link Control. Evaluation units in the third quarter. Cost: $695.	136
C&C Technology S-ARC-01 ARCnet interface	An ARCnet interface board that interfaces to STEbus. Designed around the SMC9026 VLSI controller. Acts as an intelligent slave. Operated by three on-chip single-byte registers, with additional bits for board-level control of interrupts. Cost: $858.	137
Force Computers CPU-26 SBC	A single-board computer based on Motorola's 68020 CPU and 68882 floating-point coprocessor. Exchanges messages with up to 20 other CPUs on the same backplane. Features a 32-bit DMA controller, SCSI interface, four serial ports, and VMEPROM, the company's real-time operating system kernel. Cost: $4,990 for 12.5 MHz, $5,390 for 16.7 MHz, $5,690 for 20 MHz.	138
GoldStar Technology GS230 Computer	An AT-compatible desktop computer with a 16-MHz 80286 CPU, 1 Mbyte of RAM expandable to 5 Mbytes on-board, and a 5.25-inch 1.2-Mbyte floppy drive. Comes with serial and parallel ports and a real-time clock with battery backup. Includes a socket for an 80287 math coprocessor. Comes with MS-DOS 3.3 and GW Basic. Cost: $2,195.	139
GoldStar Technology GS100 Computer	An XT-compatible computer based on a 10-MHz 8088 CPU with 640 Kbytes of RAM and a 5.25-inch 360-Kbyte floppy disk drive. Includes serial and parallel ports and a real-time clock with battery backup. Has a socket for an 8087 math coprocessor. Cost: $945.	140
Grid Systems Gridlite XL Laptop Computer	A laptop computer based on the 80C86 processor running at 8 MHz, switchable to 4.77 MHz. Comes standard with a 1.44-Mbyte 3.5-inch floppy drive, 128 Kbytes of RAM, and parallel, serial, and RGB ports. Options include a 20-Mbyte 2.5-inch hard drive, 640 Kbytes of RAM, modems, and an 8087 coprocessor. Cost: $1,950.	141
Lomas Data Products LDP-SX Motherboard	A motherboard based on the Intel 80386SX processor. Includes two serial ports configured as COM1 and COM2, a parallel port configured as LST1, an AT-compatible floppy disk controller, and a SCSI port. Cost: begins at $1,050.	142
MetraByte MV2 Frame grabber	A frame/line grabber interface board for IBM PC, XT, AT, or compatible computers. Accepts RS-170, NTSC, or CCIR compatible monochrome video input signals from interlaced, noninterlaced, or line scan cameras. Acquires the video signal through a single input channel at a speed of 30 frames/s. Digitizes to a frame resolution of 512×512×8 bits. Cost: $1,200.	143
Pacific Microcomputers PM683RT SBC	A single-board computer based on the Motorola 68030 microprocessor with optional floating-point processor. Handles interrupts from as many as 20 different sources. Comes with two serial I/O channels capable of synchronous or asynchronous operation. Optional SCSI controller. Compatible with VRTX, OS-9, and Unix. Cost: $3,975.	144
PC Craft PCC 2200/16SX Computer	A desktop computer based on the 16-MHz Intel 80386SX microprocessor. Comes with 1 Mbyte of RAM expandable to 4 Mbytes on the system board, a 1.2-Mbyte 5.25-inch floppy disk drive with a 1:1 interleave hard disk controller, two serial ports, and one parallel port. Cost: $2,100.	145
Rapid Technology DCE-400 Add-on board	An add-on board for IBM AT computers. Reduces the size of large image files for storage. Compressed images are CCITT group III or group IV compatible. Based on an integrated data compression and expansion VLSI chip. Uses double-buffered image memory. Both buffers are dual ported. Cost: $695.	146
Star Technologies VP-3 Array processor	A 32-bit CMOS array processor operating at speeds up to 150 Mflops. Based on an MIMD pipelined systems architecture. Can be attached simultaneously to three hosts, or combined with other Star processors to a single host. Comes rack-mount or stand-alone. Volume shipments at the end of 1989. Cost: $115,000.	147

9th SYMPOSIUM on COMPUTER ARITHMETIC

ARITH 9

SEPTEMBER 6-8, 1989
Santa Monica,
California

Sponsored by the Technical Committees on Computer Architecture (TCCA)
and VLSI (TCVLSI), IEEE Computer Society

in cooperation with IFIP WG 2.5 and UCLA Computer Science Department

ANNOUNCEMENT

The **Ninth Symposium on Computer Arithmetic** will be held at the Loews Santa Monica Beach Hotel in Santa Monica, California on September 6-8, 1989. The program consists of ten sessions in a single track:

Number Systems
On-Line Arithmetic (2 sessions)
Function Generation, Division, and Square Root
Adders and Convolvers
Floating-Point Arithmetic
Square Root and Division
Processors
Residue Arithmetic
Optical Arithmetic, Error Detection, and Encoding

Copies of the Advance Program are available from the Conference Registrar, Ms. Patty Patterson.

Conference Registration:

	Before August 15, 1989	After August 15, 1989
Members	$155.00	$200.00
Non-members	$200.00	$250.00
Full-time Students	$70.00	$95.00

Please make checks payable to: **ARITH-9** and send them (with your name, affiliation, address, phone number, and IEEE Member number) to:

Ms. Patty Patterson
TRW Defense Systems Group
One Space Park (R2/1094)
Redondo Beach, CA 90278 USA

Call (213) 812-0788 if you have questions.

Hotel Reservations:

Reservations should be made directly by calling the Loews Santa Monica Beach Hotel at (800) 223-0888. Request the conference rate and make your reservations by August 15, 1989 to get the special rate of $118.00 (single or double occupancy).

IEEE COMPUTER SOCIETY

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

CONFERENCES

Editor: Edmund L. Gallizzi, Computer Science Dept., Eckerd College, St. Petersburg, FL 33733; (813) 864-8272; Compmail, e.gallizzi

Symposium in Jerusalem focuses on memory, parallel architectures

Uri Weiser, Intel Israel
Gabriel M. Silberman, Carnegie Mellon University

Memory in its many manifestations within a computer system was the dominant topic of this year's program at the 16th ACM SIGArch/IEEE Computer Society Annual International Symposium on Computer Architecture in Jerusalem, May 28 to June 1. Sessions covered cache coherence and synchronization in multiprocessors, uniprocessor cache organization and evaluation, instruction fetching and caching, memory systems in vector processors, and memory in interconnection networks. Another major topic was parallel architectures and dataflow.

The technical program featured 46 papers by researchers from 10 countries. There were 300 participants, including representatives from Australia, Japan, Korea, Taiwan, Western and Eastern Europe, and a strong contingent from the US. Participants from Israeli academia and high-technology industries made up about a third of the total. A poster session, featuring 13 additional papers, attracted much interest during the first evening.

Michael Yoeli and Gabriel M. Silberman of the Technion in Haifa, Israel, acted as general cochairs, while Jean Claude Syre of ECRC in Munich, Germany, served as program chair, assisted by program vice-chairs Arvind of the Massachusetts Institute of Technology, John Gurd of the University of Manchester, UK, and Masaru Kitsuregawa of the University of Tokyo.

Preceding the symposium, a series of computer architecture workshops, sponsored by ACM SIGArch, were held in the resort city of Eilat on the Red Sea. These workshops, held for the first time in conjunction with the symposium, addressed five different topics in computer architecture, ranging from a "Status Report on Dataflow" through "Cache Architectures in Multiprocessors" and "Interconnect Architectures for Shared Memory Multiprocessors" to "Architectural Support for Declarative Languages."

More than 100 participants engaged in two days of intensive and productive discussions, followed by two days of even more intensive sunbathing, sightseeing, snorkeling, and scuba diving. Workshop minutes will appear in a forthcoming issue of *Computer Architecture News*, the SIGArch newsletter. Topic selection and coordination for the workshops were led by Trevor N. Mudge of the University of Michigan.

While workshop participants traveled from Eilat to Jerusalem via the Dead Sea and Masada, the tutorial program started on May 28. Coordinated by Daniel Tabak of George Mason University, this program attracted 127 professionals to five different topic areas. "Designing High-Yield VLSI Systems," by Israel Koren of the University of Massachusetts, Amherst; "Par-

1989 Eckert-Mauchly Award goes to Seymour Cray

Seymour Cray received the 1989 ACM/IEEE Eckert-Mauchly Award for his continuous series of outstanding contributions to the design, development, and production of large scientific computers. The award was presented during a banquet at the 16th ACM SIGArch/IEEE Computer Society Annual International Symposium on Computer Architecture, in Jerusalem. Robert Levy of Cray Research in France accepted the award on Cray's behalf and spoke briefly of Seymour Cray and his continuing commitment to building ever faster processors.

While working for Engineering Research Associates in the 1950s, Cray had design responsibility for a major portion of the ERA 1103, the first commercially successful scientific computer. During those years, he worked with the full range of computer technologies, from vacuum tubes and magnetic amplifiers to transistors.

He founded Cray Research in 1972 to design and build the world's highest performance general-purpose supercomputers. Upon its introduction in 1976, the Cray-1 established a new standard in supercomputing. The Cray-2, another step forward, was introduced in 1985. Currently he is developing a successor to the Cray-2 system and is exploring gallium arsenide technology.

Seymour Cray received the 1989 ACM/IEEE Eckert-Mauchly Award for outstanding contributions to the design, development, and production of large scientific computers.

allel Processing Systems,'' by H.J. Siegel of Purdue University; and ''RISC Systems, Architecture and Implementation,'' by Veljko M. Milutinovic of Purdue University and Daniel Tabak, program coordinator, shared the first day. The tutorial program was completed June 1, when Alan J. Smith of UC Berkeley presented ''Cache Memories'' and Lubomir Bic of UC Irvine and Jean Luc Gaudiot of the University of Southern California presented ''Dataflow versus von Neumann: Two Ends of a Spectrum.''

The symposium itself opened in a relaxed, informal (no ties) atmosphere on May 29. During the opening session Michael Yoeli was awarded a Certificate of Acknowledgement by the Israel Section of the IEEE and the Israel Chapter of the IEEE Computer Society for his pioneering activities and contribution to the development of computer science and engineering, and for his stimulating guidance and education of generations of computer scientists and engineers in Israel. Yonah Z. Lavi, chair of the Israel Chapter of the IEEE Computer Society, presented the award.

Following this presentation, John A. Darringer of IBM Research gave an opening address on ''Trends in Technology and Systems.'' Darringer reviewed current frontiers in such diverse fields as integrated circuits, fiber optics, magnetic media, and displays, and elaborated on their possible impact on computer systems. In his keynote address Michael W. Blasgen, also from IBM Research, talked about ''Evolution to Post-RISC Architecture,'' drawing an analogy between computer architecture and traditional architecture. In Blasgen's view, the post-RISC era, which we are in now, parallels that of postmodernism and is characterized by continuing efforts to achieve maximum performance for sequential machines alongside multiprocessor research and development. Furthermore, he cited out-of-order execution and multiexecution units as current directions in post-RISC architectures.

Special-purpose machines for logic simulation, graphics, and Prolog completed the technical program, together with an invited talk by Mario Tokoro of Sony Corporation and Keio University. Tokoro, who heads Sony's Computer Science Research Laboratory, spoke on ''Concurrent Objects: A Way to Efficiently Utilize Hardware Parallelism,'' insightfully reviewing various existing object-oriented languages and the suitability of their paradigms to hardware parallelism.

There were two panel sessions, one on SIMD versus MIMD architectures and the other on real-time issues in hardware/software architectures, chaired by Andre M. van Tilborg of ONR and Abraham Waksman of AFOSR, respectively. The first panel noted that most parallel machines consist of eight to 32 processors and follow the MIMD model. However, few work according to their original specifications, and many are obsolete when completed.

Proceedings from the symposium are available from the Computer Society Press, Los Alamitos, California, order number 1948. Member price is $45, nonmember price is $90. Phone (800) CS-BOOKS or (714) 821-8380.

Workshop explores diverse views of computer security

Joshua D. Guttman
The Mitre Corporation

The Second Computer Security Foundations Workshop, sponsored by the IEEE Computer Society's Technical Committee on Security and Privacy, was devoted to current approaches to a pair of problems. The first is to formally define computer security, or aspects of it; the second is to assess the usefulness of various formal notions for establishing the security of particular systems.

Three panels and 19 short presentations filled two days, June 12 and 13, at the Franconia Inn in Franconia, New Hampshire. Jon Millen of Mitre Corporation served as both general chair and program chair.

Most formal definitions of computer security policies now focus on the idea of information flow. Because older approaches tended to define mechanisms (for instance, access control mechanisms), they are usually considered too implementation dependent to serve as definitions of security policies. For example, consider the ''safety problem'' for an access control mechanism. This problem asks whether there exists a sequence of commands after which a particular process will have access to a particular computer resource, such as a file. In 1976 Harrison, Ruzzo, and Ullman demonstrated that this problem is algorithmically undecidable.

In a recent survey paper Ravi Sandhu of George Mason University described work demonstrating that the conditions under which an access control mechanism has a decidable safety problem are quite delicate. In many cases now known to be decidable, the complexity of the decision procedure is still unknown. Hence, it hardly seems appropriate to use an access control policy to specify the overall security policy that a computer system is intended to obey.

A second drawback to access control models of security, such as the Bell-LaPadula model recommended in a Department of Defense publication, is the problem of covert channels. A covert channel is a way for one process to signal information to another by manipulating system behavior, perhaps through file locks or even the CPU scheduler, which are not normally guarded by access controls.

For these two reasons the preferred strategy for specifying a security policy is to use the notion of information flow. This method is particularly suitable for modeling military security policies, sometimes called ''mandatory'' security policies. Such a policy uses a partially ordered set of security ''markings,'' or ''levels.'' Information flow policies aim to formalize the idea that data objects classified at a particular level — or processes active at that level — should be prevented from receiving information except from data objects or processes at a lower level in the ordering. That is, information always flows upward. But how is information flow to be defined?

There are essentially two routes to an answer. One is to not define the flow of information itself but to define the situations in which information does *not* flow between two objects. This is the approach taken by noninterference (Goguen and Meseguer) and a considerable body of continuing work, including McCullough's. The idea, in essence, is that information does not flow from one process (the potential sender) to another (the potential recipient) if for every history (that is, possible sequential run of the computer system) there is another history in which the recipient has the same input-output experience but the sender has made no inputs whatsoever.

The alternative approach is to adopt a view like Shannon's, which has information flowing between two objects when variation in the value of one constrains the possible values of the other. This idea, in an all-or-nothing form according to which information either does or does not flow between two objects, is called nondeducibility security, as discussed by Sutherland. However, a quantitative form derived from Shannon can also be used to analyze information flow through covert channels in a computer system, as Jon Millen demonstrated.

The quantitative approach was developed in two papers at the workshop, one by Jon Millen and the other by Todd Wittbold, both of Mitre. Wittbold, in the more detailed of the two papers, modeled a system as a deterministic finite-state machine with two users: a transmitter and a receiver. Each has a set of inputs (commands), which generate outputs dependent on the current state of the machine. A covert channel is then modeled by a *sampling policy*, which is simply a function determining the receiver's next input on the basis of his previous inputs and outputs.

Given a fixed sampling policy, the transmitter can select inputs in such a way as to transmit some class of possible messages to the receiver. The capacity of a state machine is then defined as the least upper bound of the (Shannon) capacities of all the covert channels determined by possible sampling policies. It is proved that by using codes, the machine can be made to transmit data at rates arbitrarily close to this capacity.

The paper also developed an algebraic method for computing the capacity of such a state machine. Work of this kind appears to allow a more informative analysis of the bandwidth of covert channels in secure computer systems.

A panel on "What is Information Flow Security?" moderated by Robert Morris, chief scientist at the National Computer Security Center, seemed to agree that nondeducibility and the Shannon information-theoretic techniques are preferable to their competitors — the noninterference definition and the access control models. They appear to be better grounded because they are based on a theory of the nature of information, and they also appear to offer a finer classification of security threats.

Nevertheless, the older, noninterference method was also represented in several papers, including two discussing a particular secure system development effort, Honeywell's Lock. Lock is the first case in which the noninterference approach is used directly to demonstrate the security of a computer system design.

The most controversial panel, "In Search of Foundations," was moderated by William Young of Computational Logic, with Leo Marcus (Aerospace), John McLean (Naval Research Laboratory), and Mark Nadel and Jon Millen (Mitre) participating. Young noted that there is no comprehensive, broadly accepted definition of what computer security is, and posed the challenge of finding one. He compared this task to the effort, undertaken in the 1920s and '30s, to find a formal definition of the intuitive notion of a mechanically computable function.

Of the four panelists, only Leo Marcus, and to a lesser extent John McLean, tried to meet the challenge head-on. Indeed, even Marcus's proposal was larded with reservations. There are certainly many aspects of computer security that are not formal at all — for instance, the practical issue of how to protect magnetic tapes. Moreover, Marcus argued, a formalization of security is essentially a "framework or language to state security policies." This is quite unlike the notion of recursive function, which simply defines a set of functions on natural numbers. John McLean, who also described a proposed partial definition of security, concluded by arguing that at least two different kinds of security policy are currently needed. First, a quite abstract definition based on noninterference or nondeducibility is desirable to establish a general standard. On the other hand, more concrete models, such as the Bell-LaPadula model, are more useful for designing a secure computer system.

This panel and the discussion after it led, I believe, to a more balanced appraisal of the role formal definitions play in guiding the development of secure systems. Although abstract definitions provide a rigorous account of what is achieved by a particular security policy, many issues cannot be stated at this level of generality. Hence, the choice of one abstract definition or another depends on particular examples. These examples furnish the information needed to determine what does and does not need securing in particular computer systems.

The proceedings from this workshop are available from the Computer Society Press, Los Alamitos, California, and contain the majority of the papers. The order number is 1955, and prices are $25 for members and $50 for nonmembers. Phone (800) CS-BOOKS or (714) 821-8380.

August 1989

Crypto 89, Aug. 20-24, Santa Barbara, Calif. Sponsor: Int'l Assoc. for Cryptologic Research. Contact Kevin McCurley, IBM Research, K53/802, 650 Harry Rd., San Jose, CA 95120-6099, phone (408) 927-1708.

IJCAI 89, 11th Int'l Joint Conferences on Artificial Intelligence, Aug. 20-26, Detroit. Sponsors: American Assoc. for Artificial Intelligence et al. Contact Wolfgang Bibel, Computer Science, Univ. of British Columbia, 6356 Agricultural Rd., Vancouver, B.C., V6T 1W5, Canada, phone (604) 228-3061.

Working Conf. on Engineering for Human-Computer Interaction, Aug. 21-25, Napa Valley, Calif. Sponsor: IFIP. Contact Leonard Bass, Software Engineering Inst., Carnegie Mellon Univ., Pittsburgh, PA 15213-3890; or Claus Unger, Praktische Informatick II, Univ. of Hagen, Feithstr 140, D-5800 Hagen, West Germany.

12th Int'l Congress on Cybernetics, Aug. 21-15, Namur, Belgium. Sponsor: Int'l Assoc. for Cybernetics. Contact IAC, Palais des Expositions, Place Andre Rijckmans, B-5000, Namur, Belgium, phone (32) 81-73-52-09.

28th Technical Symp.: Aug. 24, Gaithersburg, Md. Sponsors: Washington, DC, chapter of ACM and NIST. Contact Milton S. Hess, American Management Systems, 1525 Wilson Blvd., Arlington, VA 22209, (703) 841-5942.

ITC 89, Int'l Test Conf., Aug. 27-31, Washington, DC. Cosponsor: IEEE Philadelphia Section. Contact Doris Thomas, ITC 89, PO Box 264, Mt. Freedom, NJ 07970, phone (201) 895-5260.

Congress 89, 11th World Computer Congress, Aug. 28-Sept. 1, San Francisco. Sponsor: IFIP. Contact Congress 89, PO Box 18-P, Denver, CO 80218, phone (303) 831-6338; Adrian J. Basili, AT&T, 30 Knightsbridge Rd., Piscataway, NJ 08854; or Stephen Yau, Univ. of Florida, CIS Dept., Rm. 301, Gainesville, FL 32611, phone (904) 335-8006.

AIME 89, Second European Conf. on Artificial Intelligence in Medicine, Aug. 29-31, London. Cosponsors: European Society of Artificial Intelligence in Medicine, British Medical Informatics Society. Contact AIME 89, Gothic House, Barker Gate, Nottingham, England NG1 1JU.

CALENDAR

Φ In the accompanying Calendar, the IEEE Computer Society logo indicates the conferences the society is sponsoring and participating in; additional conference sponsors are also listed. Other conferences of interest to our readers are included, as well.

For inclusion in Call for Papers or Calendar, submit information **six weeks before the month of publication** (i.e., for the **October 1989** issue, send information for receipt by **August 15, 1989**) to Chuck Governale, Calendar Dept., *Computer*, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720.

September 1989

ASE 89, Int'l Conf. on Applications of Supercomputers in Engineering, Sept. 5-7, Southampton, England. Contact Liz Newman, Computational Mechanics Inst., Ashurst Lodge, Ashurst, Southampton, SO4 2AA, England, UK, phone 44 (0) 42-12-92-853.

12th Int'l Conf. on Fault-Tolerant Systems and Diagnostics, Sept. 5-7, Prague, Czechoslovakia. Sponsor: Czechoslovak Scientific and Technical Society. Contact Jan Hlavicka, Czech Technical Univ., Dept. of Computers, Prague, CSSR.

ICIP 89, Int'l Conf. on Image Processing, Sept. 5-8, Singapore. Sponsors: IEEE Singapore Section, Nat'l Univ. of Singapore. Contact Meeting Planners Pte. Ltd., 100 Beach Rd., No. 33-01, Shaw Towers, Singapore 0718, Republic of Singapore, phone (65) 297-2822.

Φ **Arith 9, Ninth Symp. on Computer Arithmetic, Sept. 6-8,** Santa Monica, Calif. Contact Algirdas Avizienis, Computer Science Dept., Univ. of California at Los Angeles, 3732 Boelter Hall, Los Angeles, CA 90024, phone (213) 825-3028.

Seventh Pacific Northwest Software Quality Conf., Sept. 10-12, Portland, Ore. Contact Dick Hamlet, Computer Science Dept., Portland State Univ., PO Box 751, Portland, OR 97207-0751.

European Software Engineering Conf., Sept. 11-15, Warwick, England. Sponsor: British Informatics Society Ltd. Contact BISL, 13 Mansfield St., London W1M 0BP, UK, phone 01-637-0471.

Midcon 89, Sept. 12-14, Rosemont, Ill. Sponsor: IEEE. Contact Midcon 89, 8110 Airport Blvd., Los Angeles, CA 90045, phone (800) 421-6816.

Φ **Canadian Conf. on Electrical and Computer Engineering, Sept. 17-20,** Montreal. Cosponsor: Canadian Society of Electrical Engineering. Contact Micha Avni, CCECE 89, PO Box 577, Desjardins Postal Station, Montreal, Que., Canada, H5B 1B7, phone (514) 283-0004; or Paul E. Allard, Transportation Development Center — Transport Canada, 200 Rene-Levesque Blvd. West, Suite 601, West Tower, Montreal, Que., Canada, phone (514) 283-0004.

Workshop on Unix System Administration, Sept. 18-19, Somerset, N.J. Contact Thomas V. Pottanat, NYNEX, 500 Westchester Ave., 2D-4, White Plains, NY 10604, phone (914) 683-2186; or Thomas M. Raleigh, Bell Communications Research, 445 South St., MRE 2A-343, Morristown, NJ 07960-1910, phone (201) 829-4321.

HCI Int'l 89, Third Int'l Conf. on Human-Computer Interaction, Sept. 18-22, Boston. Cosponsors: IEEE Systems, Man, and Cybernetics Society et al. Contact Michael J. Smith, Industrial Engineering Dept., Univ. of Wisconsin, 1513 University Ave., Madison, WI 53706, phone (608) 263-6329.

Network Management and Control Workshop, Sept. 19-21, Tarrytown, N.Y. Sponsors: IEEE et al. Contact Ivan Frisch, Center for Advanced Technology in Telecommunications, Polytechnic Univ., 333 Jay St., Brooklyn, NY 11201, phone (718) 260-3050.

10th Systems Science Int'l Conf., Sept. 19-22, Wroclaw, Poland. Contact Jerzy Swiatek, Technical Univ. of Wroclaw, Inst. of Control and Systems Engineering, Janiszewski St. 11/17, 50-370 Wroclaw, Poland.

Φ **Compsac 89, 13th Int'l Computer Software and Applications Conf., Sept. 20-22,** Kissimmee, Fla. Contact IEEE Computer Society, 1730 Massachusetts Ave. NW, Washington, DC 20036-1903, phone (202) 371-1013; or Stephen S. Yau, CIS Dept., 801 CSE, Univ. of Florida, Gainsville, FL 32611, phone (904) 335-8006.

Fifth Int'l Conf. on Image Analysis and Processing, Sept. 20-22, Positano, Italy. Sponsors: Int'l Assoc. for Pattern Recognition et al. Contact Gabriella Sanniti di Baja, c/o Istituto di Cibernetica, C.N.R., I-80072 Arco Felice, Naples, Italy, phone 39 (81) 867-1255.

Φ **ASIC 89, Application-Specific Integrated Circuit Seminar, Sept. 25-28,** Rochester, N.Y. Contact Lynne M. Engelbrecht, 170 Mt. Read Blvd., Rochester, NY 14611, phone (716) 328-2310; or Jon K. Edwards, Eastman Kodak, Dept. 434, Fl. 1, Bldg. 9, Rochester, NY 14650, phone (716) 726-9222.

Φ **Computational Intelligence 89 Symp., Sept. 25-29,** Milan, Italy. Cosponsor: ACM. Contact Luc Steels, Free Univ. Brusselles, AI Lab, Pleinlaan 2, Gebouw K 2 B-1050, Brussels, Belgium.

Int'l Electronics Manufacturing Technology Symp., Sept. 25-27, San Francisco. Sponsor: IEEE Computer Society Components, Hybrids, and Manufacturing Technology Society. Contact Diana Bendz, IBM Corp., Dept. 363/033-3, 1701 North St., Endicott, NY 13760, phone (607) 755-3521.

Third Int'l Workshop on Distributed Algorithms, Sept. 26-28, Nice, France. Contact Jean-Claude Bermond, LRI, Bat 490, Universite Paris-Sud, 91405 Orsay, France.

Φ **Performance Evaluation, Reliability, and Exploitation of Computer Systems, Sept. 26-29,** Walbrzych, Poland. Sponsors: Polish Informatic Society et al. Contact George J. Anders, Stations and Underground Section, Electrical Research Dept., Ontario Hydro, 800 Kipling Ave., KR 151, Toronto, Ont., Canada M8Z 5S4, phone (416) 231-4111.

Φ **Second IEEE Workshop on Workstation Operating Systems, Sept. 27-28,** Pacific Grove, Calif. Contact Joseph Boykin, Encore Computer, 257 Cedar Hill St., Marlborough, MA 01752, phone (508) 460-0500.

27th Allerton Conf. on Communication, Control, and Computing, Sept. 27-29, Monticello, Ill. Sponsor: Univ. of Illinois at Urbana-Champaign. Contact J.V. Medanic or P.R. Kumar, Allerton Conf., Coordinated Science Lab, Univ. of Illinois, 1101 W. Springfield Ave., Urbana, IL 61801-3082.

October 1989

Fourth Knowledge Acquisition for Knowledge-Based Systems Workshop, Oct. 1-6, Banff, Canada. Sponsor: American Assoc. for Artificial Intelligence. Contact John H. Boose, Advanced Technology Center, Boeing Computer Services, 7L-64, PO Box 24346, Seattle, WA 98124, phone (206) 865-3253.

ICCD 89, IEEE Int'l Conf. on Computer Design, Oct. 2-4, Cambridge, Mass. Contact Sumit Das Gupta, IBM, Bldg. 306, ZIP 3A1, Hopewell Junction, NY 12533, phone (914) 894-0540; or Giovanni DeMicheli, Stanford Univ., Center for Integrated Systems, Rm. 129, Stanford, CA 94305, phone (415) 725-3632.

CAPE 89, Computer Applications in Production Engineering Conf., Oct. 2-5, Tokyo. Sponsors: IFIP, IPSJ, JSPE. Contact Business Center for Academic Societies of Japan, 2-40-14, Hongo, Bunkyo-ku, Tokyo 113, Japan, phone 81 (3) 817-5831.

OOPSLA 89, Fourth Conf. on Object-Oriented Programming Systems, Languages, and Applications, Oct. 2-6, New Orleans. Sponsor: ACM. Contact OOPSLA 89, c/o Jeff McKenna, McKenna Consulting Group, 8825 SW Umatilla St., Tualatin, OR 97062.

Second Int'l Workshop on Software Configuration Management, Oct. 3-6, Princeton, N.J. Cosponsors: ACM, Gesellschaft für Informatik. Contact Thomas Murphy, Siemens Research, 755 College Rd. E, Princeton, NJ 08540, phone (609) 734-6560.

1989 IEEE Ultrasonics Symp., Oct. 3-6, Montreal. Sponsor: IEEE Ultrasonics, Ferroelectrics, and Frequency Control Society. Contact Herman van de Vaart, Allied-Signal, Box 1021R, Morristown, NJ 07960, phone (201) 455-2482; or Narendra K. Batra, Code 6385, Naval Research Lab, Washington, DC, 20375-5000, phone (202) 767-3505.

ITU-COM 89, First World Electronic Media Symp., Oct. 3-8, Palexpo, Geneva, Switzerland. Contact ITU-COM 89 Secretariat, Int'l Telecommunication Union, Place des Nations, CH-1211 Geneve 20, Switzerland, phone 41 (22) 99-5190.

Fifth Int'l Conf. on Computer-Aided Test Processing and Its Applications, Oct. 4-6, Cambridge, Mass. Cosponsors: Inst. for Numerical Computation and Analysis et al. Contact John Miller, INCA, PO Box 2, Dun Laoghaire, County Dublin, Ireland, phone 353 (1) 772-941.

Workshop on Experiences with Building Distributed and Multiprocessor Systems, Oct. 5-6, Fort Lauderdale, Fla.. Cosponsors: Usenix et al. Contact George W. Leach, PO Box 2826, MS LG-129, Largo, FL 34649-2826, phone (813) 530-2376.

Eighth Symp. on Reliable Distributed Systems, Oct. 10-12, Seattle. Contact Jane W.S. Liu, Computer Science Dept., Univ. of Illinois, 1101 W. Springfield Ave., Urbana, IL 61801-3082, phone (217) 333-6769 or 0135.

14th Conf. on Local Computer Networks, Oct. 10-12, Minneapolis. Contact Ron Rutledge, Martin Marietta Energy Systems, Oak Ridge Nat'l Lab, Bldg. 4500N, MS 6271, PO Box 2008, Oak Ridge, TN 37831-6271, phone (615) 576-7643.

Fifth Int'l Software Process Workshop, Oct. 10-13, Kennebunkport, Maine. Cosponsors: Rocky Mountain Inst. of Software Engineering et al. Contact Dewayne Perry, AT&T Bell Labs, Rm. 3D-454, 600 Mountain Ave., Murray Hill, NJ 07974, phone (201) 582-2529.

Int'l Symp. on Computer Architecture and Digital Signal Processing, Oct. 11-14, Hong Kong. Sponsor: IEEE. Contact Wan-Chi Siu, Electronic Engineering Dept., Hong Kong Polytechnic, Hung Hom, Kowloon, Hong Kong, phone (852) 3-638344, ext. 496.

Fourth Int'l High-Level Synthesis Workshop, Oct. 15-19, Kennebunkport, Maine. Cosponsor: ACM. Contact Gaetano Borriello, Computer Science Dept., FR 35, Univ. of Washington, Seattle, WA 98195, phone (206) 543-1695.

Data and Knowledge Systems for Manufacturing and Engineering, Oct. 16-18, Gaithersburg, Md. Cosponsor: ACM. Contact Lawrence A. Rowe, Computer Science Div. — EECS, Univ. of California, Berkeley, CA 94720, phone (415) 642-5117.

CSM 89, Conf. on Software Maintenance, Oct. 16-19, Miami, Fla. Cosponsors: National Inst. of Standards and Technology et al. Contact CSM 89, IEEE Computer Society, 1730 Massachusetts Ave. NW, Washington, DC 20036-1903, phone (202) 371-1013; or Wilma Osborne, NIST, Bldg. 225, Rm. B266, Gaithersburg, MD 20899, phone (301) 975-3339.

First Int'l Conf. on Artificial Neural Networks, Oct. 17-18, London. Sponsor: Institution of Electrical Engineers. Contact Conf. Services, IEE, Savoy Pl., London WC2R 0BL, UK, phone 44 (1) 24-01-871.

Eighth Int'l Conf. on Entity-Relationship Approach, Oct. 18-20, Toronto. Sponsor: ER Inst. Contact Frederick H. Lochovsky, Computer Science Dept., Univ. of Toronto, Stanford Fleming Bldg., 10 King's College Circle, Toronto, Ont. M5S 1A4, Canada, phone (416) 978-7441.

89 IEEE Int'l Workshop on Defect and Fault Tolerance in VLSI Systems, Oct. 22-24, Tampa, Fla. Contact D.L. Landis, Electrical Engineering Dept., Univ. of South Florida, Tampa, FL 33620.

IEEE Int'l Workshop on Tools for Artificial Intelligence, Oct. 23-25, Fairfax, Va. Contact Nikolas G. Bourbakis, Machine Vision Research Lab, ECE Dept., SITE, George Mason Univ., 4400 University Dr., Fairfax, VA 22030, phone (703) 425-3930.

Int'l Workshop on Defect and Fault Tolerance in VLSI Systems, Oct. 23-25, Tampa, Fla. Contact Charles H. Stapper, Dept. A23, Bldg. 861-1, IBM, Essex Junction, VT 05452, phone (802) 769-6655.

1989 Beijing Int'l Conf. on System Simulation and Scientific Computing, Oct. 23-26, Beijing, China. Sponsors: Society for Computer Simulation et al. Contact Wen Chuan-Yuan, Beijing Inst. of Aeronautics and Astronautics, Beijing, China.

Second Int'l Symp. on Artificial Intelligence, Oct. 23-27, Monterrey, Mexico. Contact Francisco J. Cantu, Inst. Tecnologico de Monterrey, ITESM Suc. Correos "J," Monterrey, N.L., Mexico 64849, phone 52 (83) 58-20-00.

23rd Asilomar Conf. on Signals, Systems, and Computers, Oct. 30-Nov. 1, Pacific Grove, Calif. Cosponsors: Naval Postgraduate School, San Jose State Univ. Contact John T. Rickard, Orincon Corp., 9363 Towne Centre Dr., San Diego, CA 92121, phone (619) 455-5530.

Third IFAC Workshop on Experience with the Management of Software Projects, Oct. 30-Nov. 1, West Lafayette, Ind. Cosponsors: Purdue Univ. et al. Contact Frederic J. Mowle, School of Electrical Engineering, Purdue Univ., West Lafayette, IN 47907, phone (317) 494-3440.

FOCS 89, 30th Foundations of Computer Science Conf., Oct. 30-Nov. 1, Research Park, N.C. Contact Christos Papadimitriou, Computer Science Dept., Univ. of California at San Diego, La Jolla, CA 92093, phone (619) 534-2086; or FOCS 89, IEEE Computer Society, 1730 Massachusetts Ave. NW, Washington, DC 20036-1903, phone (202) 371-1013.

ISCIS 4, Fourth Int'l Symp. on Computer and Information Sciences, Oct. 30-Nov. 1, Cesme, Turkey. Contact Asuman Dogac, Computer Engineering Dept., Middle East Technical Univ., 06531 Ankara, Turkey.

CIPS Edmonton 89, Oct. 30-Nov. 1, Edmonton, Alta., Canada. Sponsor: Canadian Information Processing Society. Contact Wayne A. Davis, Computing Science Dept., Univ. of Alberta, Edmonton, Alta., Canada T6G 2H1, phone (403) 492-3976.

FedCASE 89, Federal CASE Conf., Oct. 30-Nov. 2, Gaithersburg, Md. Sponsor: Nat'l Inst. of Standards and Technology. Contact Margaret H. Law or Wilma M. Osborne, NIST, Technology Bldg., Gaithersburg, MD 20899, or 2107 Vermont Ave., Landover, MD 20785.

November 1989

Ninth Symp. on Small Computers in the Arts, Nov. 3-5, Philadelphia, Pa. Cosponsor: Small Computers in the Arts Network. Contact Dick Moberg, 338 S. Quince St., Philadelphia, PA 19107, phone (215) 834-1511.

ICCAD 89, IEEE Int'l Conf. on Computer-Aided Design, Nov. 6-9, Santa Clara, Calif. Cosponsor: IEEE Circuits and Systems Society. Contact Andrezej J. Strojwas, ECE Dept., Carnegie Mellon Univ., Pittsburgh, PA 15213-3890, phone (412) 268-3530; or IEEE Computer Society, 1730 Massachusetts Ave. NW, Washington, DC 20036-1903, phone (202) 371-1013.

GOMAC 89, Government Microcircuit Applications Conf., Nov. 7-9, Orlando, Fla. Sponsors: US Dept. of Energy et al. Contact Randolph A. Reitmeyer, Electronics Technology and Devices Lab., Attn.: SLCET-I, Ft. Monmouth, NJ 07703-5000, phone (201) 544-3465.

Fifth Int'l Congress on Advances in Non-Impact Printing Technologies, Nov. 12-17, San Diego, Calif. Sponsor: Society for Imaging Science and Technology. Contact Wayne Jaeger, Tektronix, PO Box 500, MS 50-321, Beaverton, OR 97077, phone (503) 627-6714.

Prociem 89, Second Conf. on Productivity through Computer Integrated Engineering and Manufacturing, Nov. 13-15, Orlando, Fla. Sponsor: Florida High Technology and Industry Council. Contact Vince Amico, Univ. of Central Florida, College of Extended Studies, PO Box 25000, Orlando, FL 32816, phone (407) 275-2123.

IFIP Int'l Workshop on Applied Formal Methods for Correct VLSI Design, Nov. 13-16, Leuven, Belgium. Cosponsors: IFIP, Interuniversity Micro Electronics Center. Contact Luc Claesen, IMEC vzw, Kapeldreef 75, B-3030 Leuven, Belgium phone 32 (16) 281-203.

Supercomputing 89, Nov. 13-17, Reno, Nev. Cosponsor: ACM. Contact F. Ron Bailey, MS 258-5, NASA Ames Research Center, Moffett Field, CA 94035, phone (415) 694-4500.

Semiconductor Manufacturing Conf., Nov. 14-15, Phoenix. Sponsor: Society of Manufacturing Engineers. Contact Karen L. Kammerer, Technical Committee Projects Dept., SME, 1 SME Dr., PO Box 930, Dearborn, MI 48121, phone (313) 271-1500.

12th Western Educational Computing Conf., Nov. 16-17, Burlingame, Calif. Sponsor: California Educational Computing Consortium. Contact Judah Rosenwald, Extended Education, NAD 153, San Francisco State Univ., 1600 Holloway, San Francisco, CA 94132, phone (415) 338-1212.

AIDA 89, Fifth Conf. on Artificial Intelligence and Ada, Nov. 16-17, Fairfax, Va. Cosponsors: George Mason Univ., Inst. for Defense Analyses, Software Productivity Consortium. Contact AIDA 89, Computer Science Dept., George Mason Univ., 4400 University Dr., Fairfax, VA 22030, phone (703) 323-2713.

Tencon 89, IEEE Region 10 Conf., Nov. 22-24, Bombay, India. Sponsor: IEEE Bombay Section. Contact Kirit J. Sheth, Hakotronics Pvt. Ltd., Victoria Gardens, Bombay 400 027, India, phone 91 (22) 872-2888.

IEEE Workshop on Interpretation of 3D Scenes, Nov. 27-29, Austin, Texas. Contact Anil K. Jain, Computer Science Dept., Michigan State Univ., A-714 Wells Hall, East Lansing, MI 48824, phone (517) 353-5150.

December 1989

WSC 89, Winter Simulation Conf., Dec. 3-6, Washington, DC. Cosponsors: Society for Computer Simulation et al. Contact Sallie Sheppard, Office of Associate Provost, 103 Academic Bldg., Texas A&M Univ., College Station, TX 77843, phone (512) 845-3210; or Philip Heidelberger, IBM Research Div., T.J. Watson Research Center, Hawthorne, PO Box 704, Yorktown Heights, NY 10598, phone (914) 789-7156.

First Int'l Conf. on Deductive and Object-Oriented Databases, Dec. 4-6, Kyoto, Japan. Cosponsors: Information Processing Society of Japan et al. Contact Won Kim, MCC, 3500 W. Balcones Center Dr., Austin, TX 78759, phone (512) 338-3439; Jean-Marie Nicholas, ECRC Arabellastr. 17, 8000 Munich 81, FRG, phone 49 (89) 926-99-110; or Shojiro Nishio, Applied Math and Physics Dept., Faculty of Engineering, Kyoto Univ., Toyonaka, Kyoto, 606 Japan, phone 81 (75) 751-2111, ext. 5513.

Fifth Aerospace Computer Security Applications Conf., Dec. 4-8, Tucson, Ariz. Contact Marshall Abrams, Mitre Corp., 7525 Colshire Dr., McLean, VA 22102, phone (703) 883-6938.

10th Real-Time Systems Symp., Dec. 5-7, Santa Monica, Calif. Contact Al Mok, Computer Science Dept., TAY 3-140C, Univ. of Texas, Austin, TX 78712, phone (512) 471-9542.

SIGSoft 89, Third Testing, Analysis, and Verification Symp., Dec. 6-8, Key West, Fla. Cosponsor: ACM. Contact Richard A. DeMillo, Purdue Univ., Computer Science Dept., West Lafayette, IN 47907, phone (317) 494-7823.

Third Int'l Workshop on Petri Nets and Performance Models, Dec. 11-13, Kyoto, Japan. Cosponsors: Society of Instrument and Control Engineers et al. Contact Shojiro Nishio, Applied Math and Physics Dept., Faculty of Engineering, Kyoto Univ., Toyonaka, Kyoto, 606 Japan, phone 81 (75) 751-2111, ext. 5513.

Fourth SIAM Conf. on Parallel Processing for Scientific Computing, Dec. 11-13, Chicago. Sponsor: Society for Industrial and Applied Mathematics. Contact SIAM Conf. Coordinator, 117 S. 17th St., 14th Floor, Philadelphia, PA 19103-5052, phone (215) 564-2929.

Int'l Conf. on CAD/CAM and Advanced Manufacturing Technology, Dec. 19-21, Jerusalem, Israel. Cosponsor: Israel Society for CAD/CAM. Contact Lawrence R. Odess, Ortra, PO Box 50432, 61500 Tel Aviv, Israel, phone 972 (3) 971-3991 or 664-825.

Ninth Conf. on Foundations of Software Technology and Theoretical Computer Science, Dec. 19-21, Bangalore, India. Contact C.E. Veni Madhavan, Tata Research Development and Design Center, 1 Mangaldas Rd., Pune 411001, India, phone (212) 662-453.

Sixth Israeli Conf. on Artificial Intelligence and Computer Vision, Dec. 26-27, Tel Aviv. Sponsor: Information Processing Assoc. of Israel. Contact IPA, PO Box 13009, 91130 Jerusalem, Israel.

January 1990

First Int'l Symp. on Artificial Intelligence and Mathematics, Jan. 3-5, Fort Lauderdale, Fla. Contact Frederick Hoffman, Mathematics Dept., Florida Atlantic Univ., PO Box 3091, Boca Raton, FL 33431-0991, phone (407) 367-3340.

HICSS 23, 23rd Hawaii Int'l Conf. on System Sciences, Jan. 3-6, Kona, Hawaii. Cosponsor: Univ. of Hawaii. Contact Luqi, Computer Science Dept., Naval Postgraduate School, Monterey, CA 93943, phone (408) 646-2735.

Winter 1990 Usenix Tech. Conf., Jan. 22-26, Washington, DC. Contact Daniel V. Klein, Software Engineering Inst., Carnegie Mellon Univ., Pittsburgh, PA 15213-3890, phone (412) 268-7791.

IEEE Int'l Conf. on Wafer Scale Integration, Jan. 23-25, San Francisco. Contact Joe Brewer, 351 White Cedar Ln., Severna Park, MD 21146, phone (301) 765-1247.

February 1990

Sixth Int'l Conf. on Data Engineering, Feb. 5-9, Los Angeles. Contact Joseph E. Urban, College of Engineering, Univ. of

Miami, PO Box 248294, Coral Gables, FL 33124, phone (305) 284-2404.

Applications Technology Conf., Feb. 12-15, San Jose, Calif. Sponsor: American Electronics Assoc. Contact Ed Teja, American Electronics Assoc., 5201 Great American Pkwy., Santa Clara, CA 95054, phone (503) 231-9914.

Eurasip Workshop on Neural Networks, Feb. 15-17, Sesimbra, Portugal. Cosponsors: European Assoc. for Signal Processing, IEEE, Instituto de Engenharia de Sistemas e Computadores. Contact Luis B. Almeida, INESC, Apartado 10105, P-1017 Lisboa Codex, Portugal, phone 351 (1) 544-607.

CSC 90, 18th Computer Science Conf., Feb. 20-22, Washington, DC. Sponsor: ACM. Contact Barbara Kyriakakis, Computer Science Dept., George Mason Univ., Fairfax, VA 22030, phone (703) 323-2318.

IEEE Compcon Spring 90, Feb. 26-Mar. 2, San Francisco. Contact Kenichi Miura, Computational Research Dept., MS B2-7, Fujitsu America, 3055 Orchard Dr., San Jose, CA 95134-2017, phone (408) 432-1300, ext. 5408 or 5723; or IEEE Computer Society, 1730 Massachusetts Ave. NW, Washington, DC 20036-1903, phone (202) 371-1013.

Third Int'l Software for Strategic Systems Conf., Feb. 27-28, Huntsville, Ala. Cosponsors: IEEE Computer Society Huntsville Chapter et al. Contact Continuing Education Div., Univ. of Alabama in Huntsville, Tom Bevill Center 285, Huntsville, AL 35899, phone (800) 448-4035 or (205) 895-6372.

March 1990

Eighth Nat'l Conf. on Ada Technology, Mar. 5-8, Atlanta. Contact Eighth Nat'l Conf. of Ada Technology, US Army Communications — Electronics Command, Attn.: AMSEL-RD-SE-CRM (Kay Trezza), Fort Monmouth, NY 07703-5000.

CAIA 90, Sixth Conf. on Artificial Intelligence Applications, Mar. 5-9, Santa Barbara, Calif. Contact IEEE Computer Society, 1730 Massachusetts Ave. NW, Washington, DC 20036-1903, phone (202) 371-1013.

Parbase 90, Int'l Conf. on Databases, Parallel Architectures, and their Applications, Mar. 6-9, Miami Beach. Sponsors: Florida Int'l Univ. et al. Contact Parbase 90, School of Computer Science, Florida Int'l Univ., Miami, FL 33199, phone (305) 554-3429 or 3386.

EDAC 90, European Design Automation Conf., Mar. 12-15, Glasgow, Scotland. Contact Gordon Adshead, ICL, 26 Albany St., Edinburgh, EH1 3QH, UK, phone 44 (61) 223-1207.

IEEE Computer Society 1990 Int'l Conf. on Computer Languages, Mar. 12-16, New Orleans. Contact Boumediene Belkouche, Computer Science Dept., Tulane Univ., PO Box 5079, 301 Stanley Thomas Hall, New Orleans, LA 70118, phone (504) 865-5840.

1990 Symp. on Interactive 3D Graphics, Mar. 18-21, Snowbird, Utah. Sponsors: US Office of Naval Research et al. Contact Rich Riesenfeld, Univ. of Utah, Computer Science Dept., 3190 Merrill Engineering Bldg., Salt Lake City, Utah 84112, phone (801) 581-8224.

IPCCC, Ninth IEEE Int'l Phoenix Conf. on Computers and Communications, Mar. 21-23, Scottsdale, Ariz. Cosponsor: IEEE Communications Society. Contact Forouzan Golshani, Computer Science Dept., Arizona State Univ., Tempe, AZ 85287, phone (602) 965-2855.

ICSE 12, 12th Int'l Conf. on Software Engineering, Mar. 26-30, Nice, France. Cosponsors: ACM et al. Contact Francois-Regis Valette, CERT/DERI, PO Box 4026-2, Ave. Edouard Belin-31055 Toulouse, France, phone (33) 61-55-71-11; ICSE 12, AFCET, 156 Bd. Pereire, 75017 Paris, France; or IEEE Computer Society, 1730 Massachusetts Ave. NW, Washington, DC 20036-1903, phone (202) 371-1013.

1990 Int'l Conf. on Extending Database Technology, Mar. 26-30, Venice, Italy. Cosponsors: EDBT Foundation et al. Contact Michael L. Brodie, Intelligent Database Systems Dept., GTE Labs, 40 Sylvan Rd., MS 62, Waltham, MA 02254, phone (617) 466-2256.

April 1990

CHI 90, Human Factors in Computing Systems 1990, Apr. 1-5, Seattle. Sponsor: ACM. Contact Assoc. for Computing Machinery, 11 W. 42nd St., New York, NY 10036.

Disco 90, Int'l Symp. on Design and Implementation of Symbolic Computation Systems, Apr. 10-12, Capri, Italy. Contact Alfonso Miola, Dip. Informatica e Sistemistica, Via Buonarroti, 12, 00185 Roma, Italy, phone 39 (6) 731-2367.

First Int'l Conf. on Systems Integration, Apr. 23-26, Morristown, N.J. Cosponsors: New Jersey Inst. of Technology et al. Contact Peter A. Ng, Computer and Information Science Dept., New Jersey Inst. of Technology, Newark, NJ 07102, phone (201) 596-3387.

COIS 90, Conf. on Office Information Systems, Apr. 25-27, Cambridge, Mass. Cosponsor: ACM. Contact Robert B. Allen, 2A-367, Bellcore, 444 South St., Morristown, NJ 07960-1910, phone (201) 829-4280 or 4315.

May 1990

10th IEEE Symp. on Mass Storage Systems, May 6-10, Monterey, Calif. Contact Bernard T. O'Lear, NCAR, PO Box 3000, Boulder, CO 80307, phone (303) 497-1268.

CompEuro 90, IEEE Int'l Conf. on Computer Systems and Software Engineering, May 7-9, Tel Aviv. Cosponsors: IEEE, Information Processing Assoc. of Israel. Contact CompEuro 90 Conf. Secretariat, c/o ORTRA, 2 Kaufman St., PO Box 50432, Tel Aviv, 61500, Israel, phone 972 (3) 664-825.

First Conf. on Visualization in Biomedical Computing, May 22-25, Atlanta. Cosponsors: Nat'l Science Foundation et al. Contact Norberto Ezquerra, Bioengineering Center, Georgia Tech, Atlanta, GA 30332, phone (404) 894-7026 or 3964.

SIGMetrics 90, May 22-25, Boulder, Colo. Sponsor: ACM. Contact Gary J. Nutt, Univ. of Colorado, Boulder, CO 80301; or Herb Schwetman, MCC, 3500 W. Balcones Center Dr., Austin, TX 78759, phone (512) 338-3428.

20th Int'l Symp. on Multiple-Valued Logic, May 23-25, Charlotte, N.C. Cosponsors: Microelectronic Center of North Carolina et al. Contact George Epstein, Computer Science Dept., Univ. of North Carolina at Charlotte, 214 Kennedy Bldg., Charlotte, NC 28223, phone (704) 547-4566; or Carolyn F. Blalock, Office of Continuing Education, Univ. of North Carolina at Charlotte, Charlotte, NC 28223, phone (704) 547-4861.

1990 American Control Conf., May 23-25, San Diego, Calif. Sponsor: American Automatic Control Council. Contact Dagfinn Gangsaas, Boeing Advanced Systems, PO Box 3707, MS 33-12, Seattle, WA 98124-2207, phone (206) 241-4348.

17th Int'l Symp. on Computer Architecture, May 28-31, Seattle. Cosponsor: ACM. Contact Jean L. Baer or Larry Snyder, Univ. of Washington, Computer Science Dept., FR-35, Seattle, WA 98195, phone (206) 543-1695.

ICDCS 10, 10th Int'l Conf. on Distributed Computing Systems, May 28-June 1, Paris. Contact R. Popescu-Zeletin, GMD-FOKUS, Hardenbergplatz 2, D-1000 Berlin 12, West Germany, phone 49 (30) 25499-206; or Jack Stankovic, Computer and Information Science Dept., Univ. of Massachusetts, Amherst, MA 01003, phone (413) 545-0720.

11th Conf. of the Canadian Applied Mathematics Society, May 29-June 1, Halifax, N.S., Canada. Cosponsors: CAMS et al. Contact Mary Meidell, Continuing Education Dept., Technical Univ. of Nova Scotia, PO Box 1000, Halifax, NS B3J 2X4, Canada, phone (902) 429-8300, ext. 2420.

June 1990

IEEE Infocom 90, Ninth Conf. on Computer Communications, June 3-7, San Francisco. Cosponsor: IEEE Communications Society. Contact Infocom 90, IEEE Computer Society, 1730 Massachusetts Ave. NW, Washington, DC 20036-1903, phone (202) 371-1013.

Ninth Int'l Conf. on Analysis and Optimization of Systems, June 12-15, Antibes, France. Sponsor: INRIA. Contact Conf. Secretariat, INRIA, Service des Relations Exterieures, Domaine de Voluceau — BP 105, 78153 Le Chesnay Cedex, France, phone 33 (1)-39-63-5500.

10th Int'l Conf. on Pattern Recognition, June 17-21, Åtlantic City, NJ. Contact Herbert Freeman, CAIP Center, 605 Hill, Rutgers Univ., New Brunswick, NJ 08903, phone (201) 932-4208.

DAC 90, 27th ACM/IEEE Design Automation Conf., June 20-29, Orlando, Fla. Cosponsor: ACM. Contact Pat Pistilli, MP Associates, 7490 Clubhouse Rd., Suite 102, Boulder, CO 80301, phone (303) 530-4333.

Second Int'l Conf. on Software Engineering and Knowledge Engineering, June 21-23, Skokie, Ill. Sponsors: Knowledge Systems Inst., Univ. of Pittsburgh, and Inst. for Information Industries, Taiwan. Contact Shi-Kuo Chang, Computer Science Dept., Univ. of Pittsburgh, 322 Alumni Hall, Pittsburgh, PA 15260, phone (412) 624-8490.

NECC 90, 11th Nat'l Educational Computing Conf., June 25-27, Nashville, Tenn. Cosponsor: Int'l Council for Computers in Education. Contact John D. McGregor, Computer Studies Dept., Murray State Univ., Murray, KY 42071, phone (502) 762-2614.

20th Int'l Symp. on Fault Tolerant Computing, June 26-28, Newcastle upon Tyne, England. Contact Neil Speirs, Computing Lab, Univ. of Newcastle upon Tyne, Newcastle upon Tyne, NE1 7RU, UK, phone 44 (91) 232-8511.

July 1990

Second Int'l Symp. on Databases in Parallel and Distributed Systems, July 2-4, Dublin, Ireland. Cosponsor: ACM. Contact Rakesh Agrawal, AT&T Bell Labs, Rm. 3D450, 600 Mountain Ave., Murray Hill, NJ 07974, phone (201) 582-2250; or David Bell, Inst. of Informatics, Univ. of Ulster, Jordanstown, County Antrim, Northern Ireland BT370QB, phone (0232) 365-131.

ICALP 90, 17th Int'l Colloquium on Automata, Languages, and Programming, July 16-20, Coventry, England. Sponsor: Int'l Computers, Ltd. Contact Computer Science Dept., Univ. of Warwick, Coventry CV4 7AL, UK, phone 44 (203) 523-194.

DIAC 90, Directions and Implications of Advanced Computing, July 28, Boston. Sponsor: Computer Professionals for Social Responsibility. Contact Douglas Schuler, Boeing Computer Services, MS 7L-64, PO 24346, Seattle, WA 98124-0346, phone (206) 865-3226.

August 1990

UPADI 90, 21th Convention of the Pan American Federation of Engineering Societies, Aug. 19-24, Washington, DC. Cosponsors: American Assoc. of Engineering Societies, American Society of Civil Engineers. Contact UPADI 90, ASCE, 345 E. 47th St., New York, NY 10017, phone (212) 705-7218.

Coling 90, 13th Int'l Conf. on Computational Linguistics, Aug. 20-25, Helsinki, Finland. Contact Hans Karlgren, KVAL, Skeppsbron 26, S-111 30 Stockholm, Sweden, phone 46 (8) 789-6683.

October 1990

InfoJapan 90, Int'l Conf. on Information Technology, Oct. 1-5, Tokyo. Sponsor: IPSJ. Contact Haruhisa Ishida, Computer Centre, Univ. of Tokyo, 2-11-16 Yayoi, Bunkyo-ku, Tokyo 113, Japan, phone 81 (3) 818-0287.

Compsac 90, 14th Int'l Computer Software and Applications Conf., Oct. 31-Nov. 2, Chicago. Contact Stephen S. Yau, CIS Dept., 801 CSE, Univ. of Florida, Gainesville, FL 32611, phone (904) 335-8006.

November 1990

ICCC 90, 10th Int'l Conf. on Computer Communication, Nov. 5-9, New Delhi, India. Sponsor: Int'l Council on Computer Communication. Contact P.P. Gupta, CMC Ltd., 1 Ring Rd., Kilokri, New Delhi-110 014, India.

ICCAD 90, IEEE Int'l Conf. on Computer-Aided Design, Nov. 12-15 Cosponsor: IEEE Circuits and Systems Society. Contact ICCAD 90, IEEE Computer Society, 1730 Massachusetts Ave. NW, Washington, DC 20036-1903, phone (202) 371-1013.

13th Western Educational Computing Conf. Nov. 15-16, Irvine, Calif. Sponsor: California Educational Computing Consortium. Contact Oliver Seely, Jr., California State Univ. at Dominguez Hills, Chemistry, 1000 E. Victoria St., Carson, CA 90747.

Cognitiva 90, Nov. 20-23, Madrid. Sponsor: AFCET. Contact Cognitiva 90, c/o AFCET, 156 Bd Pereire 75017 Paris, France, phone 33 (01) 47-66-24-19.

You Don't Work in a Vacuum.
Neither Do We.

Software is more than coding. It's also design, testing, production, validation, maintenance. That's why *IEEE Software* covers the life cycle from conception to implementation, sharing results from practitioners and researchers.

Our departments augment this coverage with insights on standards, quality assurance, human factors, management, news, and products.

We're what a technical magazine should be: Practical. Authoritative. Lucid.

For subscription information, write *IEEE Software*, 10662 Los Vaqueros Cir., Los Alamitos, CA 90720-2578, call (714) 821-8380, or use the reader-service card.

IEEE Software

The state of the art about the state of the practice.

Have you heard about our new periodical...

Transactions on Knowledge and Data Engineering?

For information on this, or any of our nine optional periodicals, circle number 200 on the reader service card.

IEEE COMPUTER SOCIETY
Membership/Circulation Dept.
10662 Los Vaqueros Circle
Los Alamitos, CA 90720
(714) 821-8380

CALL FOR PAPERS

IEEE Trans. Parallel and Distributed Computing will begin quarterly publication in January 1990. Sample subject areas include parallel and distributed architectures, parallel and distributed software, and parallel algorithms and applications. Submit paper to Tse-yun Feng, Editor-in-Chief, *IEEE TPDC*, Computer Engineering Program, Electrical Engineering Dept., Pennsylvania State Univ., University Park, PA 16802.

IEEE Micro seeks articles for general-interest issues in 1990. Suggested topics include neural networks, artificial intelligence, special-purpose computers, optical computers and interfaces, workstations, use of microprocessors in parallel computers, VHDL design, silicon compilation, and biological computing. Tutorials are welcome on all micro-related topics. Submit manuscripts to Joe Hootman, Editor-in-Chief, EE Dept., Univ. of North Dakota, PO Box 7165, Grand Forks, ND 58202, phone (701) 777-4331.

First Int'l Symp. on Artificial Intelligence and Mathematics: Jan. 3-5, 1990, Fort Lauderdale, Fla. Submit paper to Frederick Hoffman, Mathematics Dept., Florida Atlantic Univ., PO Box 3091, Boca Raton, FL 33431-0991, phone (407) 367-3340.

ICDCS 10, 10th Int'l Conf. on Distributed Computing Systems: May 28-June 1, 1990, Paris. Submit abstract and paper to R. Popescu-Zeletin, GMD-FOKUS, Hardenbergplatz 2, D-1000 Berlin 12, West Germany, phone 49 (30) 25499-206; or Jack Stankovic, Computer and Information Science Dept., Univ. of Massachusetts, Amherst, MA 01003, phone (413) 545-0720.

CSC 90, 18th Computer Science Conf.: Feb. 20-22, 1990, Washington, DC. Sponsor: ACM. Submit paper by **Aug. 15, 1989**, to Barbara Kyriakakis, Computer Science Dept., George Mason Univ., Fairfax, VA 22030, phone (703) 323-2318.

IEEE Infocom 90, Conf. on Computer Communications: June 3-7, 1990, San Francisco. Cosponsor: IEEE Communications Society. Submit paper by **Aug. 31, 1989**, to M. Gerla, IEEE Infocom 90, Computer Science Dept., UCLA, BH 3732H, Los Angeles, CA 90024, phone (213) 825-2660.

CompEuro 90, IEEE Int'l Conf. on Computer Systems and Software Engineering: May 7-9, 1990, Tel Aviv. Cosponsors: IEEE, Information Processing Assoc. of Israel. Submit abstract by **Sept. 1, 1989**, to

Call for papers and referees for *Computer*

Computer seeks articles for inclusion in upcoming special issues.

Recent Developments in Operating Systems has been selected as the theme for the May 1990 edition. Tutorial, survey, descriptive, case-study, applications-oriented, or pedagogic manuscripts are sought.

Suggested topics include

- Distributed systems
- Multiprocessor systems
- Real-time systems
- Security and privacy
- Workstations
- File systems
- Performance evaluation
- Reliability
- Virtual memory
- Future trends

Articles must not have been previously published or currently submitted for publication elsewhere.

A 300-word abstract of the manuscript is due by **August 15, 1989**, and eight copies of the full manuscript must be submitted by **October 1, 1989**, to Joseph Boykin, Encore Computer Corp., 257 Cedar Hill St., Marlborough, MA 01752, phone (508) 460-0500, ext. 2720, Internet boykin@encore.com; UUCP encore!boykin

Authors will be notified of acceptance by **December 15, 1989**. The final version of the manuscript is due no later than **February 1, 1990**.

Persons interested in serving as referees are asked to send a note indicating their technical interests and qualifications to Boykin or to Bruce Shriver, Editor-in-Chief, *Computer*, c/o the University of Southwestern Louisiana, Drawer 42730, Lafayette, LA 70504, phone (318) 231-6000, Compmail+ b.shriver; Internet: shriver@usl.edu

Voice In Computing has been selected as the theme for the August 1990 edition. Surveys, tutorials, system or application descriptions, or case-studies are solicited in the following areas:

- Computer interface to voice communications, including call management, phone answering, and voice conferencing.
- Voice manipulation by computer, including recording, storage, and editing.
- Computer voice applications, including voice mail systems and voice response systems.
- Integrated voice/data applications, including voice annotation, voice/data conferencing, and uses of ISDN.
- Voice in computer user interfaces, including use of speech-synthesis and recognition.
- Speech synthesis or recognition systems.
- Computer architectures for supporting voice, including voice support in workstations, voice servers, and distributed voice connection control.

Articles must not have been previously published or currently submitted for publication elsewhere.

A 300-word abstract of the manuscript is due as soon as possible, and eight copies of the full manuscript must be submitted by **October 1, 1989**, to Ragui Kamel, Computing Research Laboratory, Bell-Northern Research, PO Box 3511 — Station C, Ottawa, Canada K1Y 4H7, phone (613) 763-3609, fax 1(613) 763-4222, electronic mail: Ragui@BNR,CA

Authors will be notified of acceptance by **January 1, 1990**. The final version of the manuscript is due no later than **March 1, 1990**.

Persons interested in serving as referees are asked to send a note indicating their technical interests and qualifications to Bruce Shriver, Editor-in-Chief, *Computer*, c/o the University of Southwestern Louisiana, Drawer 42730, Lafayette, LA 70504, phone (318) 231-6000, Compmail+ b.shriver; Internet: shriver@usl.edu

CompEuro 90 Conference Secretariat, c/o Ortra, 2 Kaufman St., PO Box 50432, Tel Aviv, 61500, Israel, phone 972 (3) 664-825.

Applications Technology Conf.: Feb. 12-15, 1990, San Jose, Calif. Sponsor: American Electronics Assoc. Submit abstract by **Sept. 1, 1989,** to Ed Teja, American Electronics Assoc., 5201 Great American Pkwy., Santa Clara, CA 95054, phone (503) 231-9914.

EDAC 90, European Design Automation Conf.: Mar. 12-15, 1990, Glasgow, Scotland. Submit abstract/paper by **Sept. 4, 1989**, to EDAC 90 Secretariat, CEP Consultants, 26-28 Albany St., Edinburgh, EH1 3QH, UK, phone 44 (31) 557-2478.

Int'l J. Intelligent Systems plans a special issue on temporal issues in artificial intelligence. Submit paper by **Sept. 8, 1989,** to Frank D. Anger, Computer Science Dept., Florida Inst. of Technology, Melbourne, FL 32901, phone (407) 725-9756; or Ken Ford, Computer Science Div., Univ. of West Florida, Pensacola, FL 32514, phone (904) 474-2551.

1990 American Control Conf.: May 23-25, 1990, San Diego, Calif. Sponsor: American Automatic Control Council. Submit paper by **Sept. 15, 1989**, to Eliezer Gai, C.S. Draper Laboratory, 555 Technology Square, MS 4B, Cambridge, MA 02139, phone (617) 258-2232.

1990 Symp. on Interactive 3D Graphics: Mar. 18-21, 1990, Snowbird, Utah. Sponsors: US Office of Naval Research et al. Submit extended abstract by **Sept. 15, 1989,** to Rich Riesenfeld, Univ. of Utah, Computer Science Dept., 3190 Merrill Engineering Bldg., Salt Lake City, Utah 84112, phone (801) 581-8224.

SIGMetrics 90, : May 22-25, 1990, Boulder, Colo. Sponsor: ACM. Submit paper by **Sept. 15, 1990,** to Mary K. Vernon, Computer Sciences Dept., Univ. of Wisconsin at Madison, 1210 W. Dayton St., Madison, WI 53706.

UPADI 90, 21th Convention of the Pan American Federation of Engineering Societies: Aug. 19-24, 1990, Washington, DC. Cosponsors: American Assoc. of Engineering Societies, American Society of Civil Engineers. Submit abstract by **Sept. 15, 1989,** to UPADI 90, c/o George L. De Feis, ASCE, 345 E. 47th St., New York, NY 10017, phone (212) 705-7290.

IEEE 1990 Int'l Conf. on Computer Languages, Mar. 12-16, New Orleans. Submit paper by **Sept. 22, 1989,** to Alexander L. Wolf, AT&T Bell Laboratories, MH 3C-533, 600 Mountain Ave., Murray Hill, NJ 07974, phone (201) 582-6443.

CAIA 90, Sixth Conf. on Artificial Intelligence Applications: Mar. 5-9, 1990, Santa Barbara, Calif. Submit paper by **Sept. 29, 1989,** to Se June Hong, IBM T.J. Watson Research Center, Rm. 31-206, PO Box 218, Yorktown Heights, NY 10598, phone (914) 945-2265.

10th Int'l Conf. on Pattern Recognition: June 17-21, 1990, Atlantic City, NJ. Submit paper by **Sept. 30, 1989,** to John Jarvis, Rm. 48601, AT&T Bell Labs, Holmdel, NJ 07733, phone (201) 949-2392.

11th Conf. of the Canadian Applied Mathematics Society: May 29-June 1, 1990, Halifax, N.S., Canada. Cosponsors: CAMS et al. State intent to submit abstract by **Sept. 30, 1989,** and submit abstract by **Nov. 20, 1989,** to Matlur Rahman, Applied Mathmetics Dept., Technical Univ. of Nova Scotia, PO Box 1000, Halifax, NS B3J 2X4, Canada, phone (902) 420-7724.

Ninth Int'l Conf. on Analysis and Optimization of Systems: June 12-15, 1990, Antibes,

---

Formal Methods for Software Engineering has been selected as the theme for coordinated special issues of *Computer*, *IEEE Software*, and *IEEE Transactions on Software Engineering* in September 1990.

Formal methods are design and construction methods explicitly based on well-defined mathematical formalisms. Examples include VDM, Z, box structures, traces, predicate transformers, state transition systems, axiomatic data types, and many more.

These methods promise

(1) better control over the system development process through clarity and precision of specification and then of development steps and
(2) reduced error commission and persistence through rigor, systematic review, and formal analysis.

Much progress has been made in using formal methods, developing support systems for them, and evaluating their applicability on industrially oriented problems. Applications to critical systems are appearing worldwide, and there is now some commercial interest based on advances in verifiable execution environments. Several standards groups are using formal methods, and one — VDM — is undergoing the international standards process.

For the coordinated set of articles planned in September 1990, multiple submissions of different types (for example, a case study with accompanying research description) are permitted. *Computer* will publish a survey plus tutorial, *IEEE Software* will carry application case studies, and *IEEE Transactions on Software Engineering* will publish research papers.

Submit application case studies, tutorials, and surveys to Susan Gerhart, MCC Software Technology Program, 3500 W. Balcones Dr., Austin, TX 78759, phone (512) 338-3492, fax (512) 338-3899, e-mail gerhart@mcc.com

Submit research contributions to Nancy Leveson, Information and Computer Science Dept., University of California, Irvine, CA 92717, phone (714) 856-5517, fax (714) 856-4056, e-mail nancy@ics.uci.edu

Drafts must be submitted no later than **October 1, 1989.** Reviews will be completed by **March 1, 1990**, and revisions must be completed by **May 1, 1990**.

Reviewers are sought who are willing to adhere to a tight timetable for publication of these special issues.

Fault-Tolerant Systems has been selected as the theme for the **July 1990** edition. Submittals should be either tutorial or survey in nature, describe state-of-the-art designs, or report recent developments and research findings in the following areas:

- Fault-tolerant hardware
- Fault-tolerant software
- Fault-tolerant distributed systems and networks
- Modeling and evaluation
- Fault-tolerant commercial systems
- Special-purpose designs
- Emerging areas

Articles must not have been previously published or currently submitted for publication elsewhere.

A 300-word abstract of the manuscript is due as soon as possible, and eight copies of the full manuscript must be submitted by **November 1, 1989,** to Adit D. Singh, Electrical and Computer Engineering, University of Massachusetts, Amherst, MA 01003, phone (413) 545-0188, electronic mail: singh@umaecs.edu; or Singaravel Murugesan, Information Sciences Division, NASA Ames Research Center, MS 244-4, Moffett Field, CA 94035, phone (415) 694-4233, Arpanet: murugesan@pluto.arc.nasa.gov

Authors will be notified of acceptance by **February 1, 1990**. The final version of the manuscript is due no later than **March 1, 1990**.

Persons interested in serving as referees are asked to send a note indicating their technical interests and qualifications to Bruce Shriver, Editor-in-Chief, *Computer*, c/o the University of Southwestern Louisiana, Drawer 42730, Lafayette, LA 70504, phone (318) 231-6000, Compmail+ b.shriver; Internet: shriver@usl.edu

SOFTWARE DEVELOPMENT SPECIALIST

Agouron Pharmaceuticals, Inc. has an immediate opening for a programmer/computer scientist to join an ambitious receptor-based drug design software development project.

An advanced degree in computer science or chemistry, extensive programming experience in Fortran and C, and familiarity with UNIX is essential.

Experience with computational chemistry or crystallographic software is a plus.

Superb facilities and compensation package are available.

Please apply with curriculum vitae and letters of reference to:
Pat Moses, Personnel Manager

Agouron™
Pharmaceuticals, Inc.

11025 North Torrey Pines Road
La Jolla, California 92037

Late Magazines?
No Magazines?
Membership Status Problems?
No Answers To Your Complaints?

Let your Computer Society Ombudsman cut through the red tape for you.

THE COMPUTER SOCIETY
Ombudsman
10662 Los Vaqueros Circle
Los Alamitos, CA 90720
COMPMAIL+: CS.HELP

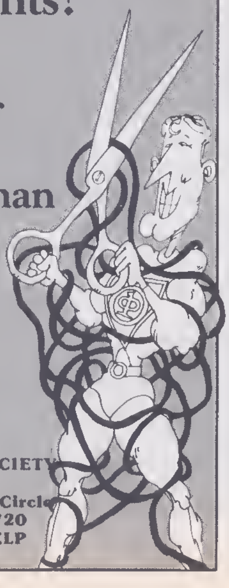

France. Sponsor: INRIA. Submit paper by **Oct. 1, 1989**, to Conference Secretariat, INRIA, Service des Relations Exterieures, Domaine de Voluceau--BP 105, 78153 Le Chesnay Cedex, France, phone 33 (1)-39-63-5500.

Disco 90, Int'l Symp. on Design and Implementation of Symbolic Computation Systems: Apr. 10-12, 1990, Capri, Italy. Submit abstract and paper by **Oct. 10, 1989**, to Alfonso Miola, Dip. Informatica e Sistemistica, Via Buonarroti, 12, 00185 Roma, Italy, phone 39 (6) 731-2367.

NECC 90, 11th Nat'l Educational Computing Conf.: June 25-27, 1990, Nashville, Tenn. Cosponsor: Int'l Council for Computers in Education. Submit paper by **Oct. 15, 1989**, to John D. McGregor, Computer Studies Dept., Murray State Univ., Murray, KY 42071, phone (502) 762-2614.

Second Int'l Symp. on Databases in Parallel and Distributed Systems: July 2-4, 1990, Dublin, Ireland. Cosponsor: ACM. Submit paper by **Oct. 16, 1989**, to Rakesh Agrawal, AT&T Bell Labs, Rm. 3D450, 600 Mountain Ave., Murray Hill, NJ 07974, phone (201) 582-2250; or David Bell, Inst. of Informatics, Univ. of Ulster, Jordanstown, County Antrim, Northern Ireland BT370QB, phone (0232) 365-131.

20th Int'l Symp. on Multiple-Valued Logic: May 23-25, 1990, Charlotte, N.C. Cosponsors: Microelectronic Center of North Carolina, Univ. of North Carolina. Submit paper by **Nov. 1, 1989**, to George Epstein, Computer Science Dept., Univ. of North Carolina at Charlotte, Charlotte, NC 28223, phone (704) 547-4566.

IEEE Trans. on Computers plans a special issue on computer arithmetic in August 1990. Submit paper by **Nov. 1, 1989**, to Earl Swartzlander, TRW Defense Systems Group, 1 Space Park, Redondo Beach, CA 90278, phone (213) 812-0791.

J. Parallel and Distributed Computing plans a special issue on software tools for parallel programming and visualization in June 1990. Submit paper by **Nov. 1, 1989**, to Lionel Ni, Computer Science Dept., Michigan State Univ., East Lansing, MI 48824-1027, phone (517) 353-4386; or K.C. Tai, Computer Science Dept., North Carolina State Univ., Raleigh, NC 27695-8206, phone (919) 737-7862.

COIS 90, Conf. on Office Information Systems: Apr. 25-27, 1990, Cambridge, Mass. Cosponsor: ACM. Submit paper by **Nov. 3, 1989**, to Fred Lochovsky, Computer Science Dept., 10 King's College Circle, Univ. of Toronto, Toronto, Ont. M5S 1A4, Canada, phone (416) 978-7441.

ICALP 90, 17th Int'l Colloquium on Automata, Languages, and Programming: July 16-20, 1990, Coventry, England. Sponsor: Int'l Computers, Ltd. Submit extended abstract or draft by **Nov. 15, 1989**, to Mike S. Paterson, Computer Science Dept., Univ. of Warwick, Coventry CV4 7AL, UK, phone 44 (203) 523-194.

17th Int'l Symp. on Computer Architecture: May 28-31, 1990, Seattle. Cosponsor: ACM. Submit paper by **Nov. 21, 1989**, to James Goodman, Computer Sciences Dept., Univ. of Wisconsin, 1210 W. Dayton, Madison, WI 53706, phone (608) 262-1204.

Coling 90, 13th Int'l Conf. on Computational Linguistics: Aug. 20-25, 1990, Helsinki, Finland. Submit paper by **Dec. 1, 1989**, to Hans Karlgren, KVAL, Skeppsbron 26, S-111 30 Stockholm, Sweden, phone 46 (8) 789-6683.

ICCC 90, 10th Int'l Conf. on Computer Communication: Nov. 5-9, 1990, New Delhi, India. Sponsor: Int'l Council on Computer Communication. Submit draft of paper by **Dec. 4, 1989**, to S. Ramani, National Centre for Software Technology, Gulmohar Cross Rd. No. 9, Juhu, Bombay-400-049, India, phone 91 (11) 620-1606.

Cognitiva 90: Nov. 20-23, 1990, Madrid. Sponsor: AFCET. Submit abstract by **Jan. 15, 1990,** to Cognitiva 90, c/o AFCET, 156 Bd Pereire 75017 Paris, France, phone 33 (01) 47-66-24-19.

InfoJapan 90, Int'l Conf. on Information Technology: Oct. 1-5, 1990, Tokyo. Sponsor: IPSJ. Submit paper by **Feb. 1, 1990,** to Haruhisa Ishida, Computer Centre, Univ. of Tokyo, 2-11-16 Yayoi, Bunkyo-ku, Tokyo 113, Japan, phone 81 (3) 818-0287.

Second Int'l Conf. on Software Engineering and Knowledge Engineering: June 21-23, 1990, Skokie, Ill. Sponsors: Knowledge Systems Inst., Univ. of Pittsburgh, and Inst. for Information Industries, Taiwan. Submit abstract and paper by **Feb. 1, 1990,** to C.Y. Hsieh, Computer Science Dept., Knowledge Systems Inst., 1153 Oak St., PO Box 576, Winnetka, IL 60093-0576.

DIAC 90, Directions and Implications of Advanced Computing: July 28, 1990, Boston. Sponsor: Computer Professionals for Social Responsibility. Submit abstract and paper by **Mar. 1, 1990**, to Douglas Schuler, Boeing Computer Services, MS 7L-64, PO 24346, Seattle, WA 98124-0346, phone (206) 865-3226.

13th Western Educational Computing Conf.: Nov. 15-16, 1990, Irvine, Calif. Sponsor: California Educational Computing Consortium. Submit paper by **Apr. 21, 1990,** to Oliver Seely, Jr., California State Univ. at Dominguez Hills, Chemistry, 1000 E. Victoria St., Carson, CA 90747.

IBM®

Be at the forefront of Computational Engineering and Science!

IBM is committed to driving the technology which will enable new developments in science and new computational problems to be solved in engineering and business.

In the heart of the Catskill region, near some of the best Eastern skiing and Hudson River boating, IBM Kingston, New York, has a team devoted to advancing this technology and its applications. To be a part of it you must be good! Some qualifications are:

- Computer Science and/or engineering skills in engineering and scientific numerical computing environments. Experience with programming of utilities and tools to support product development for MVS, VM, OS/2 and AIX (UNIX) operating systems.
- Applied math and numerical methods, linear algebra, signal processing and other mathematical library routines. Knowledge of Fortran, C and Assembler.
- Optimization, linear programming, quadratic programming.
- Application programming in UNIX environment. Knowledge of Fortran, Object Oriented Programming, C, C+ + and Assembler.
- Bridging between UNIX and non-UNIX environments. Knowledge of UNIX and one or more of the following operating systems: MVS, VM, OS/2, AIX (UNIX).
- Communication protocols, TCP/IP internals, NETBIOS, or other LAN protocols. Experience with workstations, e.g., Sun 3/XX, Sun 4/XX, Apollo, HP9000.
- Graphics, X-Window system, PHIGS, PHIGS+, GKS, PLOT3D, DECwindows, News, Display Postscript, Silicon Graphics GL, DI3000.

If this challenge excites you, you owe it to yourself to send a resume to:
IBM Corporation, Data Systems Division, Employment Department 413IE, P.O. Box 950, Poughkeepsie, NY 12602.

An equal opportunity employer

UNIX is a trademark of AT&T • X-Window is a trademark of MIT • DECwindows is a trademark of Digital Equipment Corp.
NEWS is a trademark of Sun Microsystems • AIX is a trademark of IBM • Display Postscript is a trademark of Adobe Systems
Silicon Graphics GL is a trademark of Silicon Graphics Co • DI3000 is a trademark of Precision Visuals

CAREER OPPORTUNITIES

RATES: $12.00 per line, $120 minimum charge (up to ten lines). Average five typeset words per line, eight lines per column inch. Add $10 for box number. Send copy at least one month prior to publication to: **Heidi Rex or Marian Tibayan, Classified Advertising, COMPUTER Magazine, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720; (714) 821-8380.**

THE UNIVERSITY OF TEXAS SYSTEM
Research Engineering/Scientist Associate V

The University of Texas System Center for High Performance Computing (CHPC) in Austin, Texas has an opening for a Research Engineering/Scientist Associate in the department of Applications Research and Development, working in the field of scientific visualization. The CHPC is a state supported facility offering a supercomputing capability to thirteen component campuses of the University of Texas System. Hardware currently available includes a Cray X-MP/24, a Cray X-MP /14se, and a Convex C-1, as well as a variety of support computers and workstations. This is a full-time, 12-month annual appointment with regular university benefits and competitive salary.

QUALIFICATIONS:

Doctoral degree in computer science and two years' related experience; or Master's degree in above field and six years' related experience; or Bachelor's degree in above field and eight years' related experience. Competitive applicants should have experience specializing in computer graphics and visualization techniques.

PREFERRED QUALIFICATIONS:

Experience in the use of the UNIX operating system, FORTRAN and C programming languages, and vectorization and other techniques common in the supercomputing environment. Strong motivation to develop practical applications at the leading edge of technology.

JOB DUTIES:

Close liaison with users of the Cray computers to assess visualization needs and opportunities, and the development or installation of appropriate software. Special emphasis will be given to the development of distributed graphics methods based on teamwork between the supercomputer and modern high-performance graphics workstations.

Candidates should send a letter of application, current resume, current references and salary requirements to:

Janet M. McCord
The University of Texas System
Center for High Performance Computing
Balcones Research Center
10100 N. Burnet Road
Austin, TX 78758

THE UNIVERISITY OF TEXAS IS AN EQUAL OPPORTUNITY/AFFIRMATIVE ACTION EMPLOYER.

NEW YORK UNIVERSITY
Director
Academic Computing Facility

New York University invites applications for the post of Director of the Academic Computing Facility (ACF). The ACF serves the centralized computing needs of the academic community at NYU. The ACF manages a complex of computers and network facilities and provides technical assistance to users throughout the University by means of workshops, information services and a data base archive. The ACF has a yearly budget of $4M and a staff of 62. The facilities managed by the ACF include 5 user sites with several hundred terminals and workstations, a central computing facility with several DEC and IBM mainframes and a CONVEX minisupercomputer, as well as a number of high-end Graphics workstations and peripherals. Finally, the ACF manages an ever-increasing communication network that will make the above facilities accessible to users throughout the six Manhattan centers of the University. The Director of the ACF reports to the Director of the Courant Institute of Mathematical Sciences. S/he is assisted in her/his managerial functions by an Assistant Director for Technical Services, an Assistant Director for Operations. As an ex-officio member of the ACF Policy Committee, a University-wide group composed of deans and faculty, the Director plays a leading role in advising the University in all matters relating to computer facilities, software for academic uses and networking. A qualified candidate for the position will have:

- The skills and knowledge to provide leadership, inspire cooperation, and interact fruitfully in a complicated institutional setting with a diverse user population of students and faculty from a wide variety of disciplines and professional schools.
- A critical understanding of the role of computing in the research and teaching activities of a modern university.
- Substantial technical expertise in current computer technology, in particular systems performance, networks and distributed computing.
- Several years of experience in the management of a computer-related organization, preferably in higher education.

New York University is the largest private university in the United States. Its student enrollment is over 46,000, and it has a faculty of more than 5500. Its main campus, Washington Square Center, is situated in the heart of Greenwich Village in lower Manhattan.

Applications, along with professional references, should be sent to Dr. Edmond Schonberg, Interim Chair, Search Committee for Director of ACF, New York University, 251 Mercer St., New York, NY 10012. Reviews of applications will begin in August 1989 and will continue until the position is filled. NYU encourages applications from women and members of minority groups.

SOFTWARE ENGINEER

Software Engineer for NE Ohio manufacturer to design and implement medical image processing software, specifically, computed tomography display software, using C programming language under both UNIX & VMS operating systems. Design and optimize different algorithms for image analysis and computer graphics. Work with team to integrate various software programs into final product on unique computer which runs under the PSOS (real-time multi-tasking operating system) and the REGULUS (UNIX V with real-time capabilities). No experience required in above duties but applicants will qualify with the following: M.S. in Computer Science; 3 credit hours each in Computer Graphics, Advanced Computer Graphics, Software Engineering, System Programming & Operating System & Advanced Database; 3 mos. research experience in field of parallel computation; must have used C programming language, UNIX & VMS operating system throughout 2 yr. Master Degree Program. 40 hrs/wk, 8am-5pm. $31,800/yr. Qualified applicants send resume (NO CALLS) to C. Bussard, JO# 1084447, Ohio Bureau of Employment Services, P.O. Box 1618, Columbus, OH 43216.

NEC RESEARCH INSTITUTE
Research Opportunities in Computer Sciences

The NEC Research Institute, which has just commenced operation in Princeton, New Jersey, will be conducting long term, basic research in Computer and Physical Sciences. Positions are available for senior scientists in the Computer Science division. Candidates should preferably have experience in Parallel Computing (architecture and algorithms), Software (especially languages and environments for massive parallelism), Vision, Artificial Intelligence, Theory of Computing, or Neural Networks, although the most important qualifications are prior basic research productivity, breadth of research, and future potential. Applicants should have a doctorate or equivalent qualifications and a minimum of 6 years of research experience, although more is desirable, and positions are available for very senior people. Applicants should also be US citizens or permanent residents. The Institute will provide a collegial atmosphere in which collaboration between scientists within the Institute and in other institutions will be encouraged, and publication of results in the open literature will be expected.

Competitive salaries based on experience. For further information or to apply please contact NEC Research Institute, 4 Independence Way, Princeton, New Jersey, 08540, 609-520-1555. C.W. Gear, Vice President, Computer Sciences.

Equal Opportunity Employer M/F/H/V.

UNIVERSITY OF AUCKLAND NEW ZEALAND

Two Chairs in Computer Science

The University of Auckland seeks to appoint two qualified Computer Scientists who have research and teaching skills that will enable them to make significant contributions to its rapidly developing Department of Computer Science.

The University, with over 16,000 students, holds a premier position and is sited in the heart of New Zealand's largest city. Auckland, City of Sails, is the international gateway to New Zealand, the major industrial, commercial and cultural city in the country, and offers an exceptional range of lifestyles and recreational activities.

Our Department of Computer Science, now nine years old, has 370 equivalent full-time students. In the past the department has led the way with the use of modern computing technology in teaching. Now is is poised on the brink of a second period of development following a recent review of the position of computing disciplines within the university. The new chairs will provide leadership for this new development and the associated expansion of the department's resources.

Applications are welcomed from those who believe they are qualified for these challenging positions. The successful applicants will be expected to have advanced qualifications, accomplished research records and a demonstrated history of teaching, administration and liaison with industry and commerce. The new chairs will be encouraged to foster relations with the business sector, including engagement in Consultancy activities.

The precise conditions of appointment are subject to negotiation. Further information including standard conditions of appointment and method of application are available from the Assistant Registar (Academic Appointments), University of Auckland, Private Bag, Auckland, New Zealand (FAX 64-9-799317). Applications should be forwarded by the closing date 31st AUGUST 1989.

The University of Auckland
An Equal Opportunity Employer

W B Nicoll
Registrar
University of Auckland
Private Bag
Auckland
New Zealand

COMPUTER SYSTEMS ENGINEER

Computer Systems Engineer (2 openings) for central Ohio firm to conduct research and design on graphics and database related work in a distributed telephone network traffic monitoring environment. No experience required. Must have Ph.D. in Computer Science or Physics with a graduate course in databases and operational systems. 40 hours per week (work schedule: 7:45am to 4:30 pm). Salary is $46,500 per year. Qualified applicants send resume (NO CALLS) to D. Ritchey, JO# 1197581, Ohio Bureau of Employment Services, P.O. Box 1618, Columbus, OH 43216.

UNIVERSITY OF COLORADO/BOULDER

The Department of Electrical and Computer Engineering is seeking a quality tenure-track faculty member to further enhance its research and educational activities in Computer Engineering. The successful candidate must have an outstanding academic record and significant achievement in original research as well as interest in undergraduate and graduate education. A Ph.D. degree in Electrical Engineering, Computer Engineering, Computer Science or other related field is required; salary will be commensurate with qualifications and experience. Preference will be given to candidates at the Assistant Professor level but candidates at all levels will be considered. Areas of specialization include, but are not necessarily limited to, Software Engineering, Programming Languages, Operating Systems, Hardware/Software Interface and Computer Architecture.

The University of Colorado at Boulder has a strong institutional commitment to the principle of diversity in all areas. In that spirit, we are particularly interested in receiving applications from a broad spectrum of people, including women, members of ethnic minorities and disabled individuals.

Applications for this position should be sent to: Prof. Frank S. Barnes, Dept. of Electrical and Computer Engineering, University of Colorado, Campus Box 425, Boulder, CO 80309-0425. Deadline for application is September 30, 1989.

TRINITY COLLEGE

Computer Science

Applications are invited for a tenure track position in Computer Science in the Department of Engineering and Computer Science at Trinity College starting in January 1990. The position involves undergraduate and graduate instruction and research. A doctoral degree in Computer Science or equivalent is required. Applicants will be considered in all areas of computer science that complement or strengthen our current curricular/research programs. We are interested in receiving applications from qualified women and minorities. Trinity College is an equal opportunity/affirmative action employer. Please send resume to Professor Joseph D. Bronzino, Chairman, ECS Department, Trinity College, Hartford, CT 06106. Consideration of applications will begin on May 1, 1989 and the search will remain open until the position is filled.

THE UNIVERSITY OF THE WEST INDIES, ST. AUGUSTINE, TRINIDAD

Applications are invited for the following post:

Professor of STATISTICS specializing in Applied Statistics.

ANNUAL SALARY RANGE: Professor: TT$91,788-$111,372. Pension, passages, housing, travel grant. Applications detailing qualifications and experience and naming three referees to the Registrar as soon as possible. Further particulars sent to all applicants.

UNIVERSITY OF CENTRAL FLORIDA

Post-doctoral Positions in Automated Simulation

Institute for Simulation and Training (IST)

The Institute for Simulation & Training has (2) post-doctoral positions available in the area of automatic simulator opposing force research. These positions involve the application of both object-oriented and algorithmic/functional techniques in a heterogeneous computing environment to solve problems associated with automating a network of training simulators. A working knowledge of expert systems, neural networks, data structures, and optimization techniques is strongly desired. Completed Ph.D. in Computer Science or related discipline is required. The initial appointment will be for one year, renewable for two more years depending on the availability of funds and by mutual agreement. Send resume with the names of three references to Dr. T.L. Clarke, Institute for Simulation and Training, University of Central Florida, Research Pavilion, Suite 300, 12424 Research Parkway, Orlando, FL 32826. Equal Opportunity/Affirmative Action Employer.

UNITED STATES AIR FORCE ACADEMY DEPARTMENT OF COMPUTER SCIENCE VISITING FACULTY POSITION

The Department of Computer Science of the United States Air Force Academy invites applications for a Visiting Professor/Associate Professor/Assistant Professor position. We seek qualified applicants with extensive experience teaching computer science and a record of scholarly activity. Duties will include teaching undergraduate courses, reviewing our academic program, and promoting undergraduate and faculty research. Applicant must be a U.S. citizen and should have a demonstrated commitment to undergraduate computer science programs. The appointment is normally for one year and will begin July 1990. Salary is commensurate with your current salary level. To apply, please send your vita by 1 September 1989 to: Chairman, Department of Computer Science, United States Air Force Academy, USAF Academy, CO 80840.

TULANE UNIVERSITY

The Department of Electrical Engineering at Tulane University has an immediate opening for a full-time, tenure track faculty position. Duties include teaching graduate and undergraduate courses, research, and advising students. Rank and salary are commensurate with qualifications. Requirements: Doctorate or final stages of completion of a doctorate in Electrical or Computer Engineering. Complete vitae with a minimum of three references should be sent as soon as possible to Dr. S. T. Hsieh, Department of Electrical Engineering, Tulane University, New Orleans, LA 70118-5674, (504) 865-5785. All candidates should indicate citizenship and, in the case of non-US citizens, describe their visa status. Tulane University is an equal opportunity/affirmative action employer.

COMING THIS FALL FROM THE IEEE COMPUTER SOCIETY PRESS!

Visualization in Scientific Computing

Edited by Gregory M. Nielson, *Computer Science Department, Arizona State University* and
Bruce Shriver, *Vice-President for Research, University of Southwestern Louisiana*

c. 275 pages. Fall 1989. Heavily illustrated with approximately 125 full color illustrations. Casebound, 8 1/2 X 11. ISBN 0-8186-1979-1.
Book No. 1979. Pre-Publication Prices: List $54.00/Member $42.50.

If you enjoyed the eleven visualization articles in this month's **Computer,** here is your chance to obtain them plus seven new, never before published articles in permanent hardbound form, including approximately 125 full color illustrations. The book more than doubles the visualization material in this issue -- 275 pages in all!

By drawing from many areas, some very old (such as geometry) and some relatively new (such as computer graphics), researchers are developing new and exciting tools which will allow us to more fully cooperate with the computer so that we may gain a better understanding of the world in which we live. The goal of this book is to inform the reader of these research activities and the impact they are having on how science and engineering is being done today and how it will be done in the future.

In addition to the articles in this issue of **Computer, Visualization in Scientific Computing** includes the following articles:

- "Visualization in Science and Engineering Education" by Steve Cunningham, Judith R. Brown and Mike McGrath;
- "Visualization Idioms: A Conceptual Model for Scientific Visualization Systems" by Robert B. Haber and David A. McNabb;
- "Projection Pursuit Techniques for the Visualization of High Dimensional Datasets" by Stuart L. Crawford and Thomas C. Fall;
- "A Holistic Method for Visualizing Computer Performance Measurements" by Neale Hirsch and Bruce L. Brown;
- "Graphical Analysis of Multi/Micro-tasking on Cray Multiprocessor" by Mart K. Seager, Nancy E. Werner, Mary E. Zosel and Robert E. Strout, II;
- "Computation and Manipulation of 3D Surfaces from Image Sequences" by H. Harlyn Baker;
- "Visualizing Wave Phenomenae with Seismic Rays" by V. Pereyra and J. A. Rial.

An all new Videotape on Visualization in Scientific Computing!

As a companion to the **Visualization in Scientific Computing** book, the CS Press presents a new, color, VHS videotape illustrating scientific visualization, including preprogrammed animation. Most of the chapters in the book are depicted in this video, presented by their authors. Running time: Approximately 75 minutes. List Price $49.50/Member Price $39.50

Special Discount Available!

Save 10% when you purchase the book and videotape set: Set List Price: $93.15 (a $103.50 value) Set Member Price: $73.80 (an $82.00 value)

Notes: This is advance information. Contents subject to change. Publication Date: Fall 1989. Pre-publication prices -- subject to change. Prices guaranteed for all orders mailed before September 30, 1989. Thirty-day free examination.

Order Today For 30-Day Free Examination!

Mail to:
IEEE Computer Society Press
Order Department
10662 Los Vaqueros Circle
Los Alamitos, CA 90720

Please send to me on 30 days' approval:
_____(quantity) of Visualization in Scientific Computing Book at $______ each. Subtotal: $______
_____(quantity) of Visualization Videotape at $______ each. Subtotal: $______
_____(quantity) of Book and Videotape Set at $______ per set. Subtotal: $______
Please add handling charges as indicated below. California residents add 6% sales tax

Total Enclosed: $________

Handling charges: For orders totaling $1.00 to $50.00 add $4.00; $50.01 to $75.00 add $5.00; $75.01 to $100.00 add $6.00; $100.01 add $8.00; Orders over $200.00 add $15.00.

Method of Payment:
☐Check enclosed ☐Purchase Order (number:________) ☐Visa ☐MasterCard ☐American Express

Card no.________________________ Exp. date________ Signature________________________________

Name/Title____

Organization____

Address____

City/State/Zip____

IEEE/CS Number________________________________ Phone/Telex________________________

Order toll-free!Call
1-800-CS-BOOKS
In California, call
(714)821-8380

Or fax your order!
(714)821-4010

The Institute of Electrical and Electronics Engineers, Inc. **IEEE COMPUTER SOCIETY PRESS**

BOOK REVIEWS

Editor: Guy Johnson, School of Computer Science, Rochester Institute of Technology, Rochester, NY 14623.

Illumination and Color in Computer-Generated Imagery

Roy Hall (Springer-Verlag, New York, 1989, 282 pp., $49.80)

Computer-generated graphics are finding their way into everything from video games, TV commercials, and program logos to movies, magazine, and album covers. Although many of these graphics are just flat, two-dimensional images, others are highly sophisticated, realistic, three-dimensional images requiring complex illumination and shadowing techniques.

If people working in computer graphics wanted to know which techniques add realism to computer-generated images they would have been hard pressed to find one book that contained all the information. *Illumination and Color in Computer-Generated Imagery* solves this problem by providing a collection of information on the current state of computer-graphic illumination methods. The reference book covers only this particular area and does not deal with how computer images are generated.

This is a unique book because so much of the material is normally not covered in texts or discussions on computer graphics. The first chapter examines the physics of light and the illumination process as applied to computer graphics. The next chapter gives the reader a theoretical background by introducing the basics and shortcomings of the perceptual mechanisms involved in viewing images. Hall contends that researchers cannot accurately and realistically model objects on a screen unless they understand the processes involved in illuminating, displaying, and viewing objects. After explaining the basics, Hall presents in detail the modeling and illumination techniques used to achieve realism.

Roy Hall demonstrates his knowledge of the subject matter in the first chapters. He seems to know his audience, and he keeps in mind that the book is a reference manual, not an encyclopedia. Although some of the material seems a bit disjointed, the book generally flows well and is very readable and well illustrated. Color plates are included to demonstrate the differences among the illumination models and rendering techniques.

However, Hall goes off track in the last chapter, where he tries to explain image display systems. It is evident here that Hall has strayed from his field of expertise. Certainly, some of this material needs to be covered, but much of what Hall presents does not apply to the previous material and is likely to go over the head of the intended reader. For example, Hall should have explained that a computer-generated image should be previewed on the type of equipment for which it is intended. That is, images created for a TV commercial should be viewed on normal TV monitors and compressed and color-corrected for that format. Images intended for a movie screen should be viewed on a high-quality monitor. Unfortunately, Hall doesn't get this point across very well.

This is a unique book because so much of the material is normally not covered in texts or discussions on computer graphics.

Some of the most valuable material is in the appendixes, including well-documented C language source code listings for the illumination models. Included with the programs are listings and explanations of the mathematical formulas used.

This book is an invaluable resource for those people who are seriously working and engaged in research in computer graphics. I commend Hall for collecting much disjointed material within one cover, and I hope that, in future editions, he will stick to what he knows best and stay away from what he doesn't.

Donald Pavlovich
Picker International

The Software Factory: A Fourth-Generation Software Engineering Environment

Michael W. Evans (John Wiley & Sons, New York, 1989, 308 pp., $34.95)

In his foreword to this book, Larry Yelowitz observes that the software manager of a large project faces a bewildering set of choices in establishing an environment. I hold the slightly different point of view that the manager, while disadvantaged by the lack of an industry standard project-support environment, frequently has the choice of environment made for him or her by the customer. In other words, there is frequently less choice than one might imagine.

The research and development communities of the US, Europe, and Japan are investing much of their scarcest resource, skilled software engineering professionals, in projects concerned with the specification, design, and implementation of advanced software engineering environments. Using his considerable experience of software engineering management, the author of this book writes about the type of environment that could be available in the near future. He is correctly concerned with quality and productivity and shows how engineering and management of software development can be integrated in a software factory.

I was initially puzzled by the author's use of the term "fourth-generation environment." Apparently, it refers to the type of development appropriate for applications running on distributed systems. In other words, it must provide the methods and tools to specify, design, test, and maintain networked applications. Recent years have seen explosive growth in the availability of computer-aided software engineering tools, but with few exceptions these have not been integrated into project support environments.

The book is organized into four parts. The first provides a historical perspective on computers and software engineering, discusses the need for the fourth-generation environment, and provides a broad architectural view of such an environment. The least satisfactory aspect of Part 1 is the opening chapter, which takes the form of a case study and leaves the reader in midair, uncertain as to the outcome. It is a pity the author did not return to the case study for his concluding chapter.

Part 2 stresses the importance of the engineering process within software development and deals with many data-related matters as they affect the development environment. The author pays considerable attention to configuration management. Part 3 discusses the relationship of a software engineering environment with the outside world and provides a comprehensive account of verification and validation activities. It also deals with risk categorization as it impacts on V&V and other quality control procedures.

The last part deals with automation of the software engineering environment and the problem of tailoring an environment to a project. It concludes with a section describing a work breakdown structure for a typical software project. The author comments that the structure is used by management to define and allocate work, and then incorrectly refers the reader back to Chapter 10 for further details of this activity. Unfortunately, there are many minor mistakes of this type in the book, though none are more than irritating.

The book represents a brave attempt by the author to cover an enormous amount of material.

A more serious criticism is the inconsistent level of detail. There is little or no coverage of productivity, cost estimation, and software sizing, Configuration management, on the other hand, receives comprehensive treatment. Also, apart from one reference, the author does not mention the numerous current European software engineering projects. This work should have been available to the author, since much of it has been reported in the literature and major international conferences. There are no references to Japanese experience with software factories, nor are there references to the work of the various standards organizations, though the author points out that limited standardization remains a problem. However, the good bibliography, together with the references, provides reasonable coverage of recent work in the United States.

The book represents a brave attempt by the author to cover an enormous amount of material. He does not entirely succeed, and it is difficult to imagine that the book will have much appeal to practicing software engineers. It is certainly not appropriate as a course text in software engineering project management, and, due to the uneven treatment, it may not meet the needs of project managers.

J.O. Jenkins
Imperial College

Highly Parallel Computing

George S. Almasi and Allan Gottlieb (Benjamin/Cummings Publishing, Redwood City, Calif., 519 pp., $37.95)

"Parallel processing is ready to happen," state the authors of this concept-and-overview book, which offers the basic principles of parallel computing as well as case studies of more than 20 parallel computers constructed by universities and commercial firms. All of the machines considered are designed to apply several or a great number of processors to a single problem. Large mainframes that have several processors and serve hundreds of users are viewed as machines designed for throughput and are not discussed.

The authors cite several key developments that have advanced parallel computers in performance, cost, and commercial application, including microprocessor development, switching networks on a VLSI chip that can connect several processors to numerous memory banks, the development of parallel extensions to languages like Fortran, and compilers that can find parallelism in existing programs (Parafrase). The performance levels achieved by some of these machines are truly noteworthy. For example, the IBM GF-11, a 566-microprocessor system, has attained 11 billion floating point operations per second.

The goal of parallel processing is to divide a program into sections that can be executed among separate processors. In fact, the size and number of pieces a problem is divided into (called "granularity") dictates the most suitable architecture. The authors discuss the importance of granularity in the first section. For instance, consider several large procedures that can pass data to each other; this is the idea of multiple instruction streams with multiple data streams (MIMD). At the other extreme, we can seek parallelism at the machine instruction level. Any instruction that has its operands ready will be assigned to a processor. This approach is the basis of dataflow and reduction machines.

The authors devote a chapter to showing where parallelism exists in such applications as weather code, VLSI design, database query processing, graph searching in expert and AI systems, and visual recognition systems. Once parallelism has been identified, a computational model is needed to synchronize the pieces and perform data transfer between subtasks. In the case of single-instruction, multiple-data (SIMD) machines, synchronization is not a concern because all the processors are stepping to the same instruction. In a MIMD machine, one processor may finish before the others; hence, it may have to wait for the other tasks to finish or start another job while waiting. Here the computational model provides semaphores, message techniques (such as mailboxes), and other control mechanisms to synchronize and pass data among the processors. Hardware also can play a role in synchronizing parallel programs. Memory chips have been developed that can fetch and add a number to a memory location in one operation.

The second section of the book discusses parallel software. Three chapters cover language, compilers, and operating systems. Software must define a set of subtasks, start and stop their execution, and coordinate and specify their interaction. The usual von Neumann languages: Fortran, C, Lisp, and Pascal have all been extended to accommodate parallel execution. Fortran has added statements like Parabegin and Parend. The Doall/Endall statement will iterate all the indexes used in a body of a loop. The Spawm and Mwait statements will create multiple processes and wait for each one to finish before proceeding with the rest of the program. Multilisp has been created from Lisp and features parallel expression evaluation. Concurrent Pascal offers such programming tools as processes and monitors. Even the C language has been modified with extensions similar to those found in Fortran.

Compliers also have been developed that can find parallel segments in an existing Fortran deck and generate parallel code. Parafrase, constructed by the University of Illinois, creates a dependence graph for the input code and then optimizes the program for a particular parallel computer. Generally, Parafrase has produced a speed-up factor of three or better for an input Fortran program.

Operating systems for parallel architectures are also well covered in the text. In addition to the usual functions of memory management, I/O, etc., the operating system needs to control concurrency. These systems are generally classified as master/slave or symmetric designs. In the master/slave concept, the operating system runs on one processor and assigns work to its slaves. This makes the system vulnerable to a crash if the master processor fails. The symmetric design runs parts of the operating system on any available processor. Hence, the system is not affected by the loss of a processor.

The list of parallel computers presented as case studies reads like a *Who's Who* of the parallel computer industry.

The authors have saved the best part for last. The final section covers parallel architectures and presents more than 20 parallel computers that are commercially available or exist as university projects. The machines are grouped as SIMD or MIMD computers. SIMD machines generally offer finer granularity and include vector and array processors. Processors in MIMD computers usually have their own private memory as well as shared memory; as a result, this group offers coarser granularity. The section starts with a chapter on processor-processor and processor-memory interconnections. Familiar schemes such as rings and meshes, stars, hypercubes, binary trees, and fully connected nodes are presented.

A switching network can connect processors and memory, providing dynamic connections that allow each processor to access any memory location. An interesting problem with parallel computing has been observed in these switching network designs. Often, several processors will attempt to access the same memory location to test a semaphore. This can produce a hot spot in the network as processors seek to test the same memory address. Furthermore, the commitment of paths to one location prevents other processors from obtaining paths to unrelated memory banks. The solution presented in the book is to have the network combine messages that are destined for the same location. Only one read is required and the information can be sent to all requesting processors.

The list of parallel computers presented as case studies reads like a *Who's Who* of the parallel computer industry. Under the vector processor category the following machines are presented: Cray, Cyber 205, IBM 3090, and NEC SX system. In the parallel SIMD category are Illiac IV, IBM GF-11, the Connection Machine CM-2, Goodyear C.mmp, and GE Warp. MIMD machines reviewed include the Manchester Dataflow machine, Rediflow, C.mmp, HEP (Denelcor), the NYU Ultracomputer, BBN Butterfly, IBM RP3, and the University of Illinois Cedar computer.

The authors give just enough theoretical explanation to show the trade-offs in design and programming. This book is appropriate for those readers interested in parallel computing as well as those who want a broad survey of parallel computing activity. It includes a bibliography of more than 300 papers to provide more detail on the many ideas that make up parallel computing.

Roy J. Rozman
Joliet, Illinois

ISDN — The Integrated Services Digital Network: Concepts, Methods, Systems

Peter Bocker (Springer-Verlag, New York, 1988, 250 pp., $64.50)

ISDN is a slim but very effective book dealing with almost every aspect of the evolution of public communications networks. It is intended for engineers who "design, construct, or operate communications systems" and who need an encyclopedic reference on ISDN at a summary level. Although this is a major revision of the author's German-language book on ISDN (also published by Springer-Verlag) it is very readable and gives very few hints of having been translated from another language.

ISDN's greatest strength is its fine balance between opposing viewpoints. The author discusses "true ISDN" in detail, including digital connections from subscriber to subscriber, standardized access using high-speed links, and "universal ISDN terminals" providing single-subscriber-number access to voice, data, text, and image services.

The author also discusses the existing technical hurdles facing ISDN — a world of electromechanical switches, analog transmission facilities, separate networks for voice and data, and satellite links providing high-speed connectivity at a cost in transmission latency delays. He then offers strategies for a phased introduction of ISDN and interworking with non-ISDN networks, public data networks, PBXs, and local area networks.

While the author could have presented either the user view or the network view of ISDN, he presents both perspectives as appropriate. The result is educational for both providers and consumers of ISDN services.

The meat of *ISDN* is in three chapters — "Subscriber Access," which covers the user-network interface and user signalling; "Switching in the ISDN," which concentrates on the new demands ISDN places on switching techniques and CCITT Signalling System No. 7; and "Transmission Methods in the ISDN," which discusses connectivity within the "network cloud" between subscribers.

The chapter titled "ISDN Terminals" is a good but extremely brief overview, while most of the material on what ISDN means to end users is spread throughout the book, leaving the author with nothing to say in the brief chapter titled "ISDN — the User's Viewpoint."

No book on ISDN would be complete without remarks on CCITT recommendations on ISDN and related topics, which are charted in a 13-page section and figure so heavily in the chapter references that they obscure the occasional non-CCITT reference (a minor flaw).

This book is concise yet comprehensive. The only flaw, from an engineer's point of view, is that topics worthy of book-length treatment are dealt with in paragraphs or pages. For example, the OSI reference model, about which many books have been written, is discussed in four pages, including four figures on the classic OSI protocol stack. Engineers using this book as a map into uncharted territory will be amazed at the amount of detail that lies behind this summary.

The difference between ISDN and *ISDN* is the difference between art and art appreciation, but Bocker's secondary audiences, "those engaged in teaching or research" and "a broader public with an interest in technical matters," will find the level of detail appropriate. I also highly recommend this book to engineers who are new either to the field of ISDN or to specific ISDN topics. It is clear, concise, and is excellent padding for those tackling the CCITT recommendations.

Spencer Dawkins
Bell-Northern Research

Security in Computing

Charles P. Pfleeger (Prentice Hall, Englewood Cliffs, N.J., 538 pp., $34.50)

The risks of computer security violations have grown right along with the increasing use of computers, but the technology for dealing with those risks is quite underdeveloped. Too often, the risks are only discovered after they've been exploited. The response to these problems has included increasing interest in the subject of computer security, and this textbook is intended to help solve these very real problems. The book has many problems, but despite its flaws and weaknesses, it is much better than nothing.

The book is obviously meant for university students, but the precise target audience is less obvious. Though the author teaches computer science, the book seems more suitable for nontechnical students. A few complex mathematical ideas are discussed, especially in the chapter on secure encryption systems, but the discussion is not rigorous. The author provides a few programming problems but discusses few algorithms in detail. This seems more like a book for administrators of computer systems or advanced users than for programmers.

The scope of topics is quite broad, ranging from economics — comparing the cost of security precautions to the cost of security violations and the value of data as a function of time — to software piracy and ethics. But the organization seems haphazard, related ideas reappear many times in different guises, and some of the early chapters contain substantial weaknesses that hinder understanding of the later sections. However, there is also some instructional advantage to rediscussing similar material from slightly different perspectives; this may be a strength or weakness of the book, depending on your teaching style and tempo. Still, Pfleeger doesn't develop a cohesive framework or discuss theory systematically so that students can understand what is important and what isn't. Instead, he inundates readers with ideas worth considering and events worth remembering. He also often suggests that encryption is a crucial part of the solution to a specific security problem, but the suggestions are too weak to constitute anything like a clear theme.

The book shows many signs of having been rushed to press. An especially amusing example is on p. 59, where "fi" is missing from several program segments. By the time I finished about 70 pages, I had made some very specific negative annotations. Most of these problems are just sloppy editing, but others are more basic and annoying. Many of the figures are uninformative or repetitive, and a few look like they are merely favorite lecture props, quickly and easily drawn on the blackboard but almost meaningless without the handwaving and interlocutions that go with the drawing process.

The references are coded in a most annoying fashion that is not justified by the tiny space saving. This is done even in the bibliography after each chapter, where the clarity of complete citations would be especially beneficial. Even worse, some of the references have been garbled or misplaced in the process. For example, p. 237 refers to [SAL75], which isn't decoded at all. Summaries and subject transitions are often awkward, the author's selection of chapter keywords often seems confusing and inexplicable, and the quality of the exercises is rather erratic.

Also annoying are Pfleeger's too frequent "disclaimers," meant to justify vagueness in certain examples. I agree that Pfleeger shouldn't write a cookbook for computer criminals, but neither should he use ignorance as a defense against them. Knowing precisely what security violations have occurred is the first step toward preventing their recurrence. Pfleeger's vagueness is often so extreme that, even though I recognize specific incidents to which he refers, his presentation does little to suggest appropriate countermeasures.

However, I don't want to give the impression that the book is completely displeasing and can't be recommended at all. It does have substantial weaknesses, but so does the entire field of computer security. Considering how quickly the field is evolving, Pfleeger manages to cover many current and important topics.

It would be wonderful if Pfleeger had written the definitive masterpiece on the subject, telling us how to solve all of our computer security problems. But the race between the "hackers and the crackers" will probably never end, and this book is very helpful in offering a quick overview of reasonably current standings in the race. From that perspective, the book would even be useful for self-study, though not as a reference work.

Shannon Jacobs
Austin, Texas

Current contents CS MAGAZINES

August 1989 *IEEE Micro*

THEME ARTICLES

Guest Editor's Introduction
High-Performance Microprocessors, Victor K.L. Huang

Introducing the Intel i860 64-Bit Microprocessor, Les Kohn and Neal Margulis
A million-transistor budget helps this RISC deliver balanced MIPS, Mflops, and graphics performance with no data bottlenecks.

The MC68332 Microcontroller, Joe Jelemensky et al.
This new modular microcontroller combines a high-performance CPU with a specialized time-processing unit for improved real-time control.

A Comparison of RISC Architectures, Richard S. Piepho and William S. Wu
Evaluating the newest chips for your needs can take some time and thought. Here's help in deciding what's important to consider.

SPECIAL FEATURE

A Fixed-Point DSP for Graphics Engines, Matthew Johnson
In addition to digital signal processing, the ADSP-2100 performs the kind of arithmetic required to process graphics.

August 1989 *IEEE Design & Test*

Guest Editors' Introduction
Practical BIST—Old and New Tools in Concert, Bernd Konemann and Kenneth Wagner

The Impact of Boundary Scan on Board Test, Kenneth P. Parker
While definitely not a replacement for ATE, boundary scan can still reduce test complexity, lower test cost, and increase accuracy. Those not willing to absorb the overhead of incorporating scan at the IC level may be acting silicon wise and system foolish.

Scan-Path Architecture for Pseudorandom Testing, Bulent I. Dervisoglu
This implementation of scan technology in an Apollo workstation uses a single resource to control scan and clock functions and perform pseudorandom testing of individual chips and boards.

Failure Diagnosis of Structured VLSI, John Waicukauski and Eric Lindbloom
Structured design approaches like LSSD support the efficient diagnosis of failures in VLSI chips. With a structured design, fault simulation is much faster, making it practical to simulate faults after testing.

Magazine order form

Please print or use peel-off address label

Single issue prices:
Member, $10
Nonmember, $20

Name ______

Address ______

City ______

State/ZIP/Country ______

IEEE-CS member no. (required for discount) ______

	Quantity	Price	Total
Aug. 89 *IEEE D&T*			
Aug. 89 *IEEE Micro*			

Total (payment enclosed) ______

Mail to: IEEE Computer Society Order Dept.,
10662 Los Vaqueros Circle, Los Alamitos, CA 90720 USA

For subscription information, circle number 200 on the reader service card.

NEW LITERATURE

Neural net applications. Auerbach Publishers plans to offer a new quarterly providing practical applications for neural networks. The *Journal of Neural Network Computing*, edited by Harold Szu and Judith Dayhoff, will cover tested architectures and designs. Regular columns will discuss neural net applications in image analysis and visual processing, speech analysis and recognition, knowledge processing, control systems and robotics, biological neural systems, and machine learning. The 64-page journal will cost $145 per year. Contact Auerbach Publishers, One Penn Plaza, New York, NY 10019, phone (212) 971-5000.

Tutorials tackle requirements engineering. Richard H. Thayer and Merlin Dorfman have authored *System and Software Requirements and Engineering* (Book No. 1921, 900 pp., $62 members, $90 list, casebound) and the accompanying *Standards, Guidelines, and Examples: System and Software Requirements Engineering*, (Book No. 1922, 400 pp., $41 members, $55 list, casebound). Both tutorials are scheduled for publication by the IEEE Computer Society Press in September 1989. The former emphasizes software requirements analysis and specifications, plus system engineering and its interface with software engineering. The latter combines available standards and guidelines for system and software requirements engineering. Both books ordered together cost $90 members, $130 list. Contact IEEE Computer Society Press, Order Dept., 10662 Los Vaqueros Circle, Los Alamitos, CA 90720, phone (800) 272-6657 (in California, (714) 821-8380).

Human Performance Models for Computer Aided Engineering, edited by Jerome I. Elkind et al., reviews vision and cognitive performance models; assesses their utility, validity, and integration potential; and recommends ways to integrate human factors into computer-aided engineering. The report results from NASA's request that the National Research Council study the use of task performance models in the Army-NASA Aircrew/Aircraft Integration Program. A limited number of single copies of the 309-page report are available from the Committee on Human Factors, NAS/NRC, 2101 Constitution Ave., Washington, DC 20814. When those run out, order copies from Stephan M. Zubal, National Academy Press, at the same address. The report costs $26.50 in the US and $31.75 elsewhere.

IEEE COMPUTER SOCIETY

IEEE COMPUTER SOCIETY
Membership / Subscription Application

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

BENEFITS

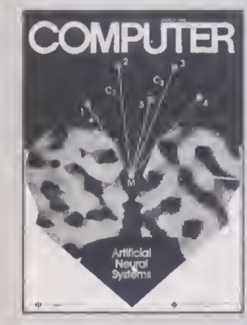

Computer
Computer comes automatically with membership. Written, reviewed, and refereed by experts, it features survey and tutorial articles covering the entire computer field, and departments such as new products, new product reviews, standards, and a reader forum called "The Open Channel." (monthly).

Technical Committees
Participate in one or more of our 33 technical committees — networks of professionals with common interests in specialty areas within computer hardware, software, and applications.

Standards Working Groups
Participate in the development of the more than 100 standards projects currently sponsored by the society in such diverse areas as software engineering, local area networks, microprocessor buses, design automation, programming languages, and standards definitions.

Computer Society Press Books
Receive discounts of up to 50% on over 600 titles covering a broad spectrum of computer science topics such as networking, communications, advanced systems, image processing, security, artificial intelligence, and design automation. Over 60 new titles are published annually.

Conferences and Tutorials
Choose from more than 100 conferences annually, ranging from large industry-oriented conferences replete with exhibits to small, highly interactive workshops. Members receive special low rates.

Schedule of Fees

To join: see item 1, 2, or 3.
To subscribe: see item 4.

Membership dues and periodical subscriptions are annualized to, and expire on, December 31. Choose full or half-year rate schedules depending on date of receipt by the Computer Society as indicated below.

		Half Year Mar 1-Aug 31	Full Year Sept 1-Feb 28
1 I don't belong to the IEEE and I want to join just the Computer Society		☐ $19.50	☐ $39.00
2 I don't belong to the IEEE and I want to join both the Computer Society and the IEEE* (Total amount to be paid includes annual dues, and regional assessment, if any.)			
I reside in Region 1-6 (United States)		☐ $44.00	☐ $88.00
I reside in Region 7 (Canada)		☐ $41.00	☐ $82.00
I reside in Region 8 (Europe, Africa, or the Middle East)		☐ $41.50	☐ $83.00
I reside in Region 9 (Latin America)		☐ $36.00	☐ $72.00
I reside in Region 10 (Asia and Pacific)		☐ $37.00	☐ $74.00
3 I already belong to the IEEE and I want to join the Computer Society. IEEE Member Number ______________		☐ $ 7.50	☐ $15.00
4 OPTIONAL PERIODICALS for new or current members	issues per year		
IEEE Computer Graphics and Applications (3061)	6	☐ $ 9.50	☐ $19.00
IEEE Design and Test (3111)	6	☐ $10.00	☐ $20.00
IEEE Expert (3151)	4	☐ $ 6.00	☐ $12.00
IEEE Micro (3071)	6	☐ $ 9.00	☐ $18.00
IEEE Software (3121)	6	☐ $ 9.00	☐ $18.00
Transactions on Computers (1161)	12	☐ $10.00	☐ $20.00
New Spring '89 *Transactions on Knowledge and Data Engineering (1471)*	4	☐ $ 5.00	☐ $10.00
Transactions on Pattern Analysis and Machine Intelligence (1351)	12	☐ $10.00	☐ $20.00
Transactions on Software Engineering (1171)	12	☐ $10.00	☐ $20.00

*ACM members who join both IEEE and the Computer Society may deduct $5 off the full-year rate; $2.50 off the half-year rate.

Total amount remitted with this application $__________

☐ Checks are accepted in Belgian, British, German, Swiss, Japanese, or U.S. currencies. U.S. checks must be drawn on a U.S. bank.

☐ Visa ☐ Master Card ☐ American Express ☐ Eurocard ☐ Diners Club

Charge Card Number ☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐

Exp. Date Mo. ☐☐ Yr. ☐☐

I hereby make application for Computer Society membership and, if elected, will be governed by IEEE's and the society's constitutions, bylaws, and statements of policies and procedures.

Full signature __________ Date __________

MAILING ADDRESS

First name __________ Middle initial(s) __________ Last name __________

Street address __________ City __________ State/Country __________ Zip __________

EDUCATION (highest level completed) Course __________ Degrees received __________ Date __________

Name of educational institution __________ Date of birth __________ ☐ Male ☐ Female

OCCUPATION Title or Position __________

Firm name __________

Firm address __________

City __________ State/Country __________ Zip __________

REFERENCE (an IEEE member; if unknown, a managerial person who knows you professionally)

Name (print in full) __________ IEEE Member No. (if applicable) __________

Street address __________

City __________ State/Country __________ Zip __________

Return with your remittance to: IEEE Computer Society, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720–2578 USA
Residents of Europe mail to: IEEE Computer Society, 13, Avenue de l'Aquilon, B-1200, Brussels, BELGIUM
Asia/Pacific residents mail to: IEEE Computer Society, Ooshima Building, 2-19-1 Minami-Aoyama, Minato-ku, Tokyo 107 JAPAN

PC0889

IEEE COMPUTER SOCIETY

A member society of the Institute of Electrical and Electronics Engineers, Inc.

Executive Committee

President: Kenneth R. Anderson*
Siemens Research & Technology
755 College Road East
Princeton, NJ 08540
(609) 734-6550

President-Elect: Helen M. Wood*
Past President: Edward A. Parrish, Jr.*

Vice Presidents

Conferences and Tutorials: Joseph E. Urban (1st VP)*
Technical Activities: Laurel V. Kaleda (2nd VP)*
Area Activities: Ned Kornfield†
Education: Gerald L. Engel†
Membership and Information: Barry W. Johnson†
Press Activities: Duncan H. Lawrie*
Publications: Sallie V. Sheppard*
Standards: Paul L. Borrill*
Secretary: Michael Evangelist*
Treasurer: Charles B. Silio†
Division V Director: Harriett Rigas†
Division VIII Director: Roy L. Russo†
Executive Director: T. Michael Elliott†

*voting member of the Board of Governors
†nonvoting member of the Board of Governors

Board of Governors

Term Expiring 1989:

Bill D. Carroll, Lansing (Chip) Hatfield,
Duncan H. Lawrie, David Pessel,
Susan L. Rosenbaum, Sallie V. Sheppard, Bruce Shriver,
Harold S. Stone, Akihiko Yamada, Marshall C. Yovits

Term Expiring 1990:

Vishwani Agrawal, Mario R. Barbacci,
Ming T. (Mike) Liu, Yale N. Patt, Donald E. Thomas,
Benjamin W. Wah, Ronald Waxman

Term Expiring 1991:

P. Bruce Berra, Paul L. Borrill, Michael Evangelist,
Ted Lewis, Raymond E. Miller,
Earl E. Swartzlander, Jr., Thomas W. Williams

Next Board Meeting

November 17, 1989, 8:30 a.m.
Bally's Hotel, Reno, NV

Senior Staff

Executive Director: T. Michael Elliott
Editor and Publisher: H. True Seaborn
Director, Computer Society Press: Eugene M. Falken
Director, Conferences and Tutorials: Anne Marie Kelly
Director, Finance and Administration: Tod S. Heisler
Director, Board and Administrative Services: Violet S. Doan

Computer Society Offices

Headquarters Office

1730 Massachusetts Ave. NW
Washington, DC 20036-1903
Phone (202) 371-0101
Telex: 7108250437 IEEE COMPSO
Fax: (202)728-9614

Publications Office

10662 Los Vaqueros Cir.
Los Alamitos, CA 90720-2578
Membership and General Information: (714) 821-8380
Publication Orders: (800) 272-6657
Fax: (714) 821-4010

European Office

13, Ave. de L'Aquilon
B-1200 Brussels, Belgium
Phone: 32 (2) 770-21-98
Fax: 32 (2) 770-85-05

Asian Office

Ooshima Building
2-19-1 Minami-Aoyama, Minato-ku
Tokyo 107, Japan
Phone: 81 (3) 408-3118
Fax: 81 (3) 408-3553

Use the Reader Service Card to obtain information on:

- Membership application—student #203, others #202
- Perodicals subscription form for individuals #200
- Periodicals subscription form for organizations #199
- Publications catalog #201
- Standards working groups list #195
- Compmail+ international electronic mail/database brochure #194
- Technical committee list/application #197
- Chapters lists, start-up procedures—student/regular #193
- Student scholarship information #192
- Awards description/nomination forms #198
- Volunteer leaders/staff directory #196
- IEEE senior member application #204

Purpose

The IEEE Computer Society advances the theory and practice of computer science and engineering, promotes the exchange of technical information among 100,000 members worldwide, and provides a wide range of services to members and nonmembers.

Membership

Members receive the acclaimed monthly magazine *Computer*, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others seriously interested in the computer field.

Publications and Activities

Computer. An authoritative, easy-to-read magazine containing tutorial and in-depth articles on topics across the computer field, plus news, conferences, calendar, interviews, and new products.

Periodicals. The society publishes six magazines and four research transactions. Refer to membership application or request information as noted above.

Conference Proceedings, Tutorial Texts, Standard Documents. The Computer Society Press publishes more than 100 titles every year.

Standards Working Groups. Over 100 of these groups produce IEEE standards used throughout the industrial world.

Technical Committees. More than 30 TCs publish newsletters, provide interaction with peers in specialty areas, and directly influence standards, conferences, and education.

Conferences/Education. The society holds about 100 conferences each year and sponsors many educational activities, including computing science accreditation.

Chapters. Regular and student chapters worldwide provide the opportunity to interact with colleagues, hear technical experts, and serve the local professional community.

European Office

Payments for Computer Society membership and publication orders are accepted by checks in Belgian, British, German, Swiss, or US currency. Checks in US funds must be drawn on a US bank. Payment may also be made by American Express, Eurocard, MasterCard, or Visa credit cards.

Asian Office

Payments for Computer Society membership and publication orders are accepted by checks in Japanese or US currency. Checks in US funds must be drawn on a US bank. Payment may also be made by electronic fund transfer to the Bank of Tokyo, Akasaka Branch, Toza acct. 0767956; the credit receiver is the IEEE Computer Society Headquarters Office. Payment may also be made by American Express, Eurocard, MasterCard, or Visa credit cards.

Ombudsman

Members experiencing problems — magazine delivery, membership status, or unresolved complaints — may write to the ombudsman at the Publications Office.

NEW FOR NETWORK ANALYSTS

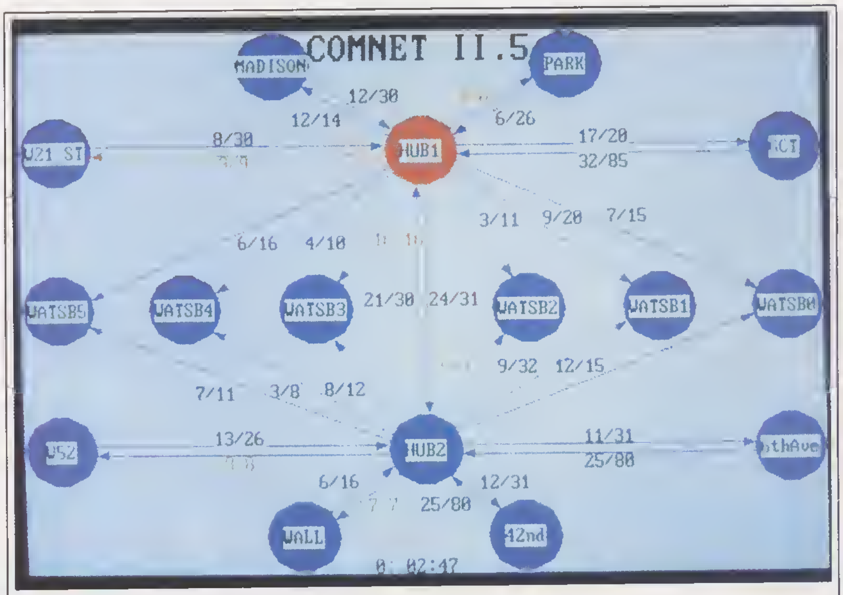

Realistic simulation of your voice or data network--quick results, no programming

COMNET II.5 now predicts the performance of your SNA, DECNET, X.25, ISDN or other packet, message, or circuit switching network

Free trial and, if you act now, free training

COMNET II.5 uses simulation to predict your network performance. You simply describe your network, traffic load, and routing algorithms.

Animated simulation follows immediately--no programming.

Easy-to-understand results

You get an animated picture of your network. Routing choices and changing levels of network utilization are apparent.

Your reports show response times, blocking probabilities, call queueing and packet delays, network throughput, circuit group utilization, and circuit group queue statistics.

Computers with COMNET II.5

COMNET II.5™ is available for most PC's, Workstations, and Mainframes.

Your network simulated

You can analyze any computer or communication network which uses circuit, message, or packet switching. Virtual-circuit or datagram operation can be modeled.

Alternate-routing and adaptive shortest-path routing algorithms are built-in.

Seeing your proposed network animated increases everyone's understanding of its operation and builds confidence in your results.

Immediate free trial information

The free trial contains everything you need to try COMNET II.5 on your computer. If you act now we will include free training.

Call Kelly Cox at (619) 457-9681, FAX (619) 457-1184. In Europe, call Richard Eve on (01) 528-7980, FAX (01) 528-7988.

Free trial offer

Free trial-see for yourself how COMNET II.5 quickly answers communication network performance questions.

Limited offer-Act now for free training.

☐ Send information on your Special University Offer.

Name

Organization

Address

City State Zip

Telephone Computer

Return to:
CACI Products Company
3344 North Torrey Pines Court
La Jolla, California 92037
Call Kelly Cox at (619) 457-9681.
FAX (619) 457-1184.

IEEE COMP

In Europe:
CACI Products Division
Regent House, 89 Kingsway
London WC2B 6RH, United Kingdom
Call Richard Eve on (01) 528-7980.
FAX (01) 528-7988.

COMNET II.5 is a trademark and service mark of CACI, INC.
©1989 CACI, INC.